

DTIC FILE COPY

RADC-TR-90-46
Final Technical Report
April 1990


AD-A223 142



3D DISPLAYS FOR BATTLE MANAGEMENT

The MITRE Corporation

H. Veron, D.A. Southard, J.R. Leger, J.L. Conway



APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

Rome Air Development Center
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700

99 06 20 041

This report has been reviewed by the RADC Public Affairs Division (PA) and is releasable to the National Technical Information Services (NTIS) At NTIS it will be releasable to the general public, including foreign nations.

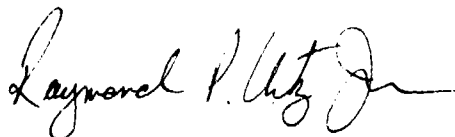
RADC-TR-90-46 has been reviewed and approved for publication.

APPROVED:



RICHARD T. SLAVINSKI
Project Engineer

APPROVED:



RAYMOND P. URTZ, JR.
Technical Director
Directorate of Command & Control

FOR THE COMMANDER:



IGOR G. PLONISCH
Directorate of Plans & Programs

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (COE) Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE

Form Approved
OPM No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Project, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE April 1990	3. REPORT TYPE AND DATES COVERED Final Oct 88 - Oct 89	
4. TITLE AND SUBTITLE 3D DISPLAYS FOR BATTLE MANAGEMENT			5. FUNDING NUMBERS C - F19628-89-C-0001 PE - 62702F PR - MOIE TA - 79 WU - 70	
6. AUTHOR(S) H. Veron, D.A. Southard, J.R. Leger, J.L. Conway				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The MITRE Corporation Bedford MA 01730			8. PERFORMING ORGANIZATION REPORT NUMBER MTR 10689	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Electronic Systems Division (TM) Hanscom AFB MA 01731-5000 Rome Air Development Center (COE) Griffiss AFB NY 13441-5700			10. SPONSORING/MONITORING AGENCY REPORT NUMBER RADC-TR-90-46	
11. SUPPLEMENTARY NOTES RADC Project Engineer: Richard Slavinski /COE/(315) 330-4031				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Stereoscopic displays enhance the ability of the viewer to comprehend spatial relationships. With the commercial introduction of flicker-free liquid crystal stereoscopic shutters (LCSS), it has become practical to incorporate stereoscopic 3D into moderately-priced graphics workstations, greatly extending the applicability of this technique. This paper reports on the results of an evaluation of stereoscopic 3D for Air Force applications. This work focused on the identification and development of optimized stereoscopic display techniques for the visualization of terrain data. Advances were made in the following areas: <u>Display Measurements.</u> Interocular crosstalk was identified as the underlying display metric that characterizes the LCSS performance and crosstalk was measured on the two commercially available LCSSs. Tests revealed some relatively minor differences in				
14. SUBJECT TERMS 3D Display			15. NUMBER OF PAGES 80	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	

UNCLASSIFIED

Block 13 (Cont'd)

performance between these two products.

Stereoscopic Projection Geometry. Existing stereoscopic drawing techniques, which have roots in the cinematography industry, were found to be inadequate for providing comfortable stereoscopic out-the-window terrain scenes, when viewed from a 19-inch diagonal CRT monitor. A stereoscopic perspective algorithm, called the "stereoscopic window" projection was developed which ensures comfortable and undistorted viewing of stereoscopic imagery.

Scene Complexity Management Algorithms. Terrain elevation and cultural feature data bases contain sampled data several orders of magnitude greater than can be represented dynamically by even the fastest graphics workstations. Algorithms and procedures were developed to manage and filter this data and construct a geometric model of the terrain that retains the information relevant to mission rehearsal applications.

Test Bed Implementation. A test bed was developed to demonstrate and evaluate the utility of various stereoscopic display techniques. A useful operating point for image complexity was determined and the scene update rate achievable with current graphics workstations was measured. Various methods of representing real cultural features, such as forests and bodies of water, and point features, such as buildings and powerline pylons were developed with practical guidelines for maximizing the effectiveness of the presentation.



Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/ _____	
Availability Codes	
Dist	Avail and/or Special
A-1	

UNCLASSIFIED

EXECUTIVE SUMMARY

Stereoscopic displays have become important tools in the fields of molecular modeling, architectural design, computer-aided design, and medicine. These displays significantly enhance the ability of the viewer to comprehend spatial relationships. With the commercial introduction of flicker-free liquid crystal stereoscopic shutters (LCSS), it has become practical to incorporate stereoscopic 3D into moderately-priced graphics workstations, greatly extending the applicability of this display technique.)

This Project 7970, "Three-Dimensional Displays for Battle Management," evaluated the 3D stereoscopic display medium for US Air Force applications. We focused on the identification and development of optimized stereoscopic display techniques for the visualization of terrain data. We applied these display techniques to a simulation of a mission planning and rehearsal system to demonstrate the utility of stereoscopic displays.

We made advances in the following areas:

- 1) Display Measurements: We identified interocular crosstalk as the underlying display metric that characterizes the LCSS performance and measured crosstalk on the two commercially available LCSSs. Our tests revealed some relatively minor differences in performance between these two products.
- 2) Stereoscopic Projection Geometry: Existing stereoscopic drawing techniques, which have roots in the cinematography industry, were found to be inadequate for providing comfortable stereoscopic out-the-window terrain scenes when viewed on a 19-inch diagonal CRT monitor. We developed a stereoscopic perspective algorithm, called the "stereoscopic window" projection, that ensures comfortable and undistorted viewing of stereoscopic imagery.
- 3) Scene Complexity Management Algorithms: Terrain elevation and cultural feature databases contain sampled data several orders of magnitude greater than can be represented dynamically by even the fastest graphics workstations. We developed algorithms and procedures to manage and filter this data, constructing a geometric model of the terrain that retains the information relevant to mission rehearsal applications.

4) Testbed Implementation. We assembled a testbed to demonstrate and evaluate the utility of various stereoscopic display techniques. We determined a useful operating point for image complexity and measured the scene update rate achievable with current graphics workstations. We experimented with various methods of representing areal cultural features, such as forests and bodies of water, and point features, such as buildings and powerline pylons, and developed some practical techniques for maximizing the effectiveness of the presentation.

The results from this effort are sufficiently novel that technical papers and presentations are being prepared for submission to several professional societies. For the next fiscal year, we plan to extend the testbed, improving both the quality of the presentation and the scene update rate, and to develop concrete applications for these stereoscopic display techniques within the Air Force.

TABLE OF CONTENTS

SECTION	PAGE
1 Introduction	1
1.1 Scope	1
1.2 Document Organization	2
2 Liquid Crystal Stereoscopic Shutter	3
2.1 Functional Description of the LCSS	3
2.2 Optical Measurements	7
2.2.1 Comments on Measurements	11
2.3 Assessment	12
3 Stereoscopic Picture Generation	13
3.1 Visual Depth Cues	13
3.2 Viewing Considerations	16
3.3 Current Methods for Stereoscopic Projection	19
3.3.1 Crossed Axes Projection Method	19
3.3.2 Parallel Axes Projection Method	19
3.3.3 Shortcomings of Current Methods	22
3.4 Stereoscopic Window Projection	24
3.4.1 Description	24
3.4.2 Discussion	24
4 Geometric Modeling of Terrain	27
4.1 Processing of Digital Terrain Elevation Data	27
4.1.1 Description of Digital Terrain Elevation Data	27
4.1.2 Terrain Feature Extraction	30
4.1.3 Laplacian Operator	31
4.1.4 Adaptive Filtering	31
4.2 Processing of Digital Feature Analysis Data	32
4.2.1 Description of Digital Feature Analysis Data	32
4.2.2 Features Database Management	34
4.2.3 SMC Image and Icon Instancing File Construction	35
4.2.4 SMC Map File and Edge Map File Creation	37

SECTION	PAGE
4.3 Merge of DTED and DFAD Data	37
4.3.1 Generating the Combined Elevation/SMC Map	39
4.3.2 Subdivision	39
4.3.3 Thresholding	39
4.4 Triangulation	40
4.4.1 Triangulated Irregular Network Model	40
4.4.2 Delaunay Triangulation	40
4.4.3 Triangular Meshes	42
4.5 Icons	42
4.5.1 Representation of DFAD Point Features Using Icons	44
4.5.2 Intervisibility Contours	47
4.6 Implementation	47
5 Terrain Scene Generation	51
5.1 Rendering	51
5.2 Complexity Mangement	54
6 Mission Planning and Rehearsal Simulation	59
6.1 Rationale	59
6.2 Simulation Description	59
6.3 Initial Reaction	62
7 Summary	63
List of References	65
Distribution List	67

LIST OF FIGURES

FIGURE		PAGE
2.1	Components of the Liquid Crystal Stereoscopic Shutter 3D Display System	5
2.2	Stereoscopic Vision Produced by the LCSS Display	6
3.1	Retinal Disparity	15
3.2	Accommodation and Convergence	18
3.3	Crossed Axes Projection	20
3.4	Parallel Axes Projection	21
3.5	Viewing a Parallel Axes Projection Stereogram	23
3.6	Viewing a Stereoscopic Window Projection Stereogram	25
4.1	Transforming DMA Data	28
4.2	Terrain Analysis	29
4.3	Feature Analysis	33
4.4	Example of 2D Surface Material Map	36
4.5	Terrain Modeling	38
4.6	Triangulation of Merged Data	41
4.7	Comparison of Threshold Techniques for DTED Model	43
4.8	Icons	45
4.9	Icons Developed for Mission Planning	46
4.10	Level of Detail	48
4.11	Intervisibility Contour	49
5.1	Interactive Terrain Rendering	52

FIGURE		PAGE
5.2	Out-The-Window View of Terrain Scene	53
5.3	Field of View Culling	56
6.1	2D Plan View Presentation of Route Planning	61

LIST OF TABLES

TABLE		PAGE
1	Interocular Crosstalk Ratios for the Tektronix LCSS	8
2	Interocular Crosstalk Ratios for the StereoGraphics Z-Screen	9
3	Parallax	16
4	Number of Samples in DTED Quadrangles	27

SECTION 1

INTRODUCTION

Recent advances in liquid crystal light valve technology have made a new and exciting display medium available to users of graphics workstations. The new technology, called the liquid crystal stereoscopic shutter (LCSS), provides the user with a dynamic, high-quality, stereoscopic view of three-dimensional (3D) imagery. Immediate advantage of this technology has been taken by members of the computer-aided design and computer-aided manufacturing (CAD/CAM), architectural, medical, and molecular modeling communities. For these groups, knowledge of the spatial relationships between objects and interaction with these objects in a 3D environment increases comprehension and work performance. When used appropriately with C³I display systems, (busy two-dimensional (2D) plan view displays), 3D stereoscopic imagery will aid in the decision-making process. MOIE (Mission Oriented Investigation and Experimentation) project 7970 entitled, "Three-Dimensional Displays for Battle Management" sets the stage to exploit the LCSS display medium for the C³I community.

Our effort uses the LCSS 3D display technology to obtain interactive out-the-window stereoscopic views of complex terrain scenes for the mission planning activities of the tactical Air Force. Stereoscopic displays for mission planning require judicious choices of system parameters to achieve good results. These parameters include selection and correct use of appropriate 3D display hardware; the development of appropriate geometric projection techniques, which provide for comfortable viewing of the resultant out-the-window stereoscopic image on a 19-inch diagonal CRT monitor; generation of a model for the complex terrain scene; and, rendering software that provides an interactive environment for the mission planner. This report presents the results of an investigation into methods for visualizing complex terrain data, which is one of the best uses of 3D stereoscopic displays.

1.1 SCOPE

This report documents the activities performed during Fiscal Year 89 on Mission Oriented Investigation and Experimentation (MOIE) project 7970, "Three-Dimensional Displays for Battle Management". These include a hardware evaluation, software development, and a final demonstration of a mission planning and rehearsal simulation on a stereoscopic workstation environment. The report is intended for technically oriented readers

although detailed knowledge of 3D graphics display techniques is not required. References 1 and 2 will provide the reader with additional background material on 3D graphics display techniques.

1.2 DOCUMENT ORGANIZATION

Section 2 evaluates available liquid crystal stereoscopic shutters.

Section 3 describes the required elements of computer-generated stereograms and includes the "stereoscopic window" projection technique that provides for comfortable viewing of stereoscopic imagery on a 19-inch cathode ray tube (CRT).

Section 4 shows how to construct a geometric model of a terrain surface from Defense Mapping Agency (DMA) databases.

Section 5 discusses the interactive terrain rendering software that provides high quality stereoscopic imagery.

Section 6 describes the application of the interactive high quality stereoscopic imagery for route planning purposes.

Section 7 summarizes the results, and describes future work.

SECTION 2

THE LIQUID CRYSTAL STEREOSCOPIC SHUTTER

In normal human vision each eye receives a separate and distinct view of the same scene. The brain processes the two views so that we see one image with depth. The complex psychophysiological interaction that gives rise to the depth sensation is called *stereopsis*. There are many devices that can provide two views of the same object on a flat display medium. These are called stereograms. With the aid of an appropriate optical device for viewing the stereogram, the observer can perceive an image in three dimensions (stereoscopic 3D). One such device is the liquid crystal stereoscopic shutter (LCSS). The LCSS is a transparent medium that is mounted on the bezel of a CRT monitor so that it overlays the surface of the CRT. The computer displays two sequentially-timed optically disparate views of the same scene. The LCSS is synchronized to the video output of the computer. When the LCSS is used with polarizing glasses, it functions as an optical switch that transmits only the left image of the stereogram to the left eye while blocking the left image to the right eye. It then switches and performs the same function for the right eye. This provides flicker-free color stereograms.

The LCSS is a relatively new device on the market. Display parameters characterizing its performance have not yet appeared in vendor data sheets. The success of the time-switched stereoscopic display is related to how well the appropriate left and right eye images are presented to their respective eyes. The image for each eye must be well isolated from the other for the observer to be able to combine these two planar images into a single stereoscopic view. The phenomenon of "bleed through" occurs when the left eye also sees the right image, or vice versa. Bleed through is also referred to as optical crosstalk between the two images. The metric that characterizes this phenomena is called the interocular crosstalk ratio. In order to evaluate crosstalk, we measured this parameter on the only two commercially available LCSSs, a Tektronix Model SGS 620, and a StereoGraphics Z-Screen.

2.1 FUNCTIONAL DESCRIPTION OF THE LCSS

Before we describe the actual measurements, we will briefly look at how the LCSS works, and how the viewer perceives a 3D stereoscopic image produced with this device.

Figure 2.1 shows a time-multiplexed LCSS display system and its critical components. The liquid crystal module mounts on the bezel of the monitor. It consists of a liquid crystal medium sandwiched between two pieces of conductively coated optical glass. A synchronization signal from the graphics computer effectively encodes each of the eye views with uniquely polarized light. By convention, the left eye view is left circularly polarized and the right eye view is right circularly polarized. The shutter system is completed with a pair of glasses that are worn by the viewer. The glasses are electrically passive and consist of oppositely polarized material which decodes the correct view for each eye. This is true for both brands of LCSS.

The liquid crystal cell remains isotropic (behaves like ordinary glass) when an electric field is impressed on the cell. The combined effect of the linear polarizer, the liquid crystal in the isotropic state, and the quarter wave plate, produces right circularly polarized light. The polarizing glasses will only pass the right image. The left lens blocks the light because it is oppositely polarized. When the electric field is removed, the liquid cell switches from an isotropic state to an anisotropic (birefringent) state. The optical configuration now produces left circularly polarized light. The polarizing glasses will now only pass the left circularly polarized light.

The polarizing glasses and the switching action of the liquid crystal shutter allow the viewer to see a time sequential stereogram. Upon receiving the proper image to each eye, the brain fuses the combination into one image with depth. Liquid crystal technology has progressed to the state where the shutter switches (changes its optical state) at the video frame rates used in state-of-the-art graphics workstations. The LCSSs evaluated in this report can switch each left and right eye view at 120 Hz so that the stereoscopic image is displayed flicker-free at 60 Hz.

The geometrical model that illustrates how an observer perceives depth from a stereogram generated by the LCSS is shown in figure 2.2. For example, let us consider a stereogram with disparate views displayed on the surface of a CRT. The disparate views of the stereogram on the display surface overlap but are horizontally displaced a finite distance from one another. This displacement is called parallax. At locations where the parallax is zero, the resultant image at that location appears to be at the surface of the screen. The stereoscopic image can also be made to appear in front of the screen. In this case, the left and right images are interchanged and a 3D stereoscopic image with negative parallax is produced. (See section 3 for more details on viewing geometry.)

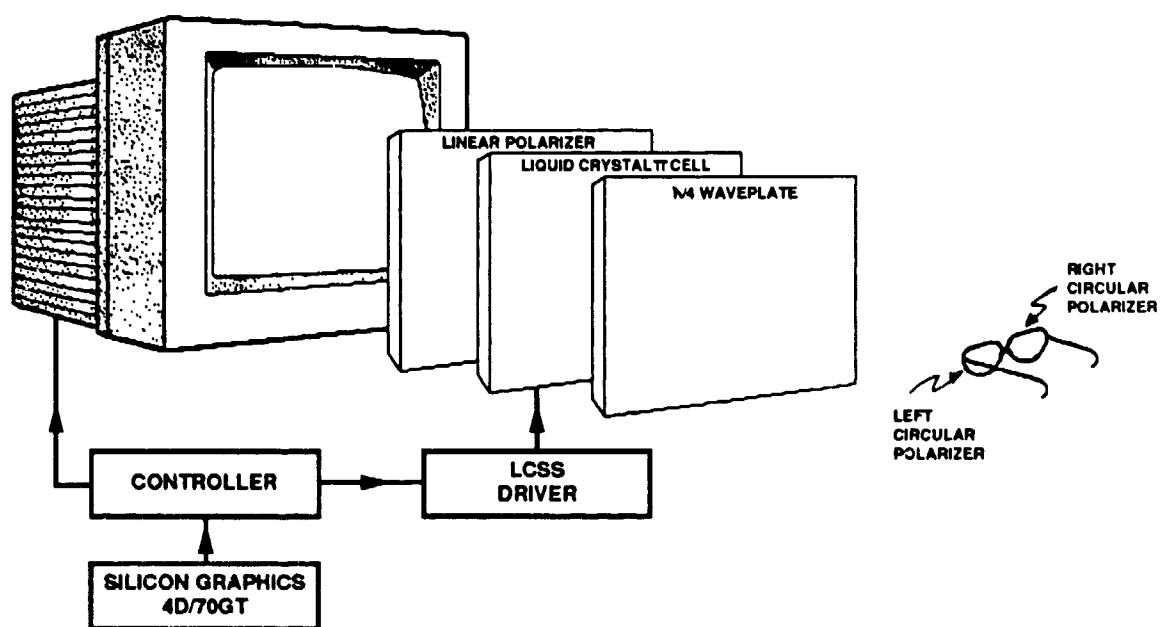


Figure 2.1. Components of the Liquid Crystal Stereoscopic Shutter 3D Display System




Figure 2.2. Stereoscopic Vision Produced by the LCSS Display

The parallax between left and right images of the stereogram on the CRT surface puts a severe requirement on the behavior of the optical combination provided by the LCSS and polarizing glasses worn by the viewer. There should be no optical crosstalk between the two images. If bleed through does appear, the ability of the viewer to fuse the resultant image may be compromised. This concern provides a rationale for quantifying the crosstalk ratio of the available LCSS devices.

2.2 OPTICAL MEASUREMENTS

The video output from the computer drives both LCSSs. This readily allowed for a side-by-side qualitative and quantitative visual assessment. We measured the transmission characteristics and the optical crosstalk between the left and right images for each system. We used a Tektronix Digital Photometer (Model J16) equipped with a one degree Narrow Angle Luminance Probe (Model J6523) to measure the average transmitted luminance from each liquid crystal shutter system.

Both liquid crystal shutter systems, within experimental error, exhibit the same transmission of about 36 percent for white light. By inserting the polarizing glasses between the LCSS and the photometer in the transmissive (ON) state, a 64 percent transmission was measured for the polarizing glasses. This implies that overall, about 24 percent of the light from the CRT reaches the observer. The luminance levels reaching the observer are between 2 to 3 footlamberts (fL) for operational CRT luminance levels between 10 to 15 fL. These levels appeared to be satisfactory for observers viewing the 3D stereoscopic imagery produced by the liquid crystal systems in a reduced room ambient.

The optical parameter that determines the performance of the LCSS stereoscopic viewing system is the interocular crosstalk. The interocular crosstalk is a ratio of the light intensity from the correct video field that gets through one of the eyeglass lenses, to the amount of light from the incorrect video field that gets through that lens. To measure interocular crosstalk, a flat field for left eye viewing only is generated. The luminance values are measured through the left and right lenses of the passive glasses. Two values are obtained. One is called LEFT ON, and the other is called RIGHT OFF. These values are indicated in tables 1 and 2. Similarly for the right eye viewing, RIGHT ON and LEFT OFF values are obtained. From these luminance values, interocular cross talk ratios LEFT ON:LEFT OFF, and RIGHT ON:RIGHT OFF, are obtained for the configurations shown in tables 1 and 2. Luminance measurements were made for each of the three primaries, red, blue, and green, and at the display screen locations indicated in tables 1 and 2. The ratios were determined for a viewer looking normal to, and then 45 degrees to, the CRT surface.

Table 1. Interocular Crosstalk Ratios for the Tektronix LCSS

Color	LEFT		RIGHT		<u>LEFT ON</u>		<u>RIGHT ON</u>		
	ON	OFF	ON	OFF	LEFT	OFF	RIGHT	OFF	
White	3.07	0.50	3.20	0.47	6.1		6.8		<div>X</div> <div>Normal</div>
Red	0.85	0.13	0.87	0.13	6.5		6.7		
Green	1.88	0.33	1.97	0.31	5.7		6.4		
Blue	0.43	0.08	0.45	0.07	5.4		6.4		
White	2.51	0.85	2.08	0.49	2.9		4.2		<div>X</div> <div>45°</div>
Red	0.68	0.18	0.60	0.14	3.8		4.3		
Green	1.54	0.56	1.28	0.31	2.8		4.1		
Blue	0.36	0.17	0.27	0.07	2.1		3.8		
White	3.56	0.69	3.25	0.46	5.2		7.1		<div>X</div> <div>Normal</div>
Red	1.03	0.17	0.90	0.11	6.1		8.2		
Green	2.18	0.46	2.02	0.31	4.7		6.5		
Blue	0.46	0.11	0.45	0.08	4.2		5.6		
White	2.69	1.01	2.28	0.43	2.7		5.3		<div>X</div> <div>45°</div>
Red	0.74	0.21	0.64	0.12	3.5		5.3		
Green	1.66	0.67	1.42	0.29	2.5		4.9		
Blue	0.37	0.19	0.30	0.07	1.9		4.3		
White	3.56	0.55	3.34	0.55	6.5		6.1		<div></div> <div>X</div> <div>Normal</div>
Red	0.98	0.14	0.91	0.15	7.0		6.1		
Green	2.21	0.36	2.08	0.36	6.1		5.8		
Blue	0.49	0.09	0.45	0.08	5.4		5.6		
White	2.43	0.86	2.33	0.56	2.8		4.2		<div></div> <div>X</div> <div>45°</div>
Red	0.66	0.18	0.66	0.16	3.7		4.1		
Green	1.52	0.56	1.45	0.36	2.7		4.0		
Blue	0.34	0.16	0.30	0.08	2.1		3.8		

Table 2. Interocular Crosstalk Ratios for the StereoGraphics Z-Screen

Color	LEFT		RIGHT		LEFT ON	RIGHT ON	
	ON	OFF	ON	OFF	LEFT OFF	RIGHT OFF	
White	2.24	0.66	2.26	0.64	3.4	3.5	<div>X</div> <div>Normal</div>
Red	0.60	0.17	0.59	0.16	3.5	3.7	
Green	1.56	0.48	1.56	0.46	3.3	3.4	
Blue	0.21	0.08	0.21	0.08	2.6	2.6	
White	1.64	0.95	1.39	0.57	1.7	2.4	<div>X</div> <div>45°</div>
Red	0.43	0.22	0.38	0.14	2.0	2.7	
Green	1.15	0.69	0.96	0.41	1.7	2.3	
Blue	0.16	0.11	0.13	0.07	1.5	1.9	
White	2.59	0.46	2.46	0.43	5.6	5.7	<div>X</div> <div>Normal</div>
Red	.72	0.10	0.68	0.09	7.2	7.6	
Green	1.84	0.36	1.74	0.33	5.1	5.3	
Blue	0.23	0.05	0.21	0.05	4.6	4.2	
White	2.13	0.93	1.50	0.30	2.3	5.0	<div>X</div> <div>45°</div>
Red	0.58	0.22	0.43	0.06	2.6	7.2	
Blue	1.53	0.70	1.06	0.24	2.2	4.4	
Green	0.18	0.09	0.12	0.04	2.0	3.0	
White	1.95	0.37	1.95	0.37	5.3	5.3	<div>X</div> <div>Normal</div>
Red	0.57	0.07	0.57	0.08	8.1	7.1	
Green	1.38	0.30	1.38	0.30	4.6	4.6	
Blue	0.17	0.05	0.17	0.05	3.4	3.4	
White	1.64	0.95	1.03	0.25	1.7	4.1	<div>X</div> <div>45°</div>
Red	0.47	0.23	0.31	0.05	2.0	6.2	
Green	1.18	0.72	0.73	0.20	1.6	3.7	
Blue	0.15	0.11	0.09	0.04	1.4	2.3	

The values of the luminance data for white light are consistent with the transmission measurements. Two to three fL reach the viewer for an ON state. The luminance values for each of the primary components, red, blue and green, add to a value of luminance approximating that of white, obtained for each of the measured ON and OFF states. This serves as a limited check on the accuracy of the photometer for both high and low values of luminance corresponding to the respective ON and OFF states.

Tables 1 and 2 show that the calculated crosstalk values for the Tektronix and Z-Screen LCSSs are different. For viewing white light normal to the CRT, the crosstalk ratio is larger for the Tektronix LCSS than for the Z-Screen LCSS. The reasons for the differences are not specifically known, but we suspect that the split screen structure of the Tektronix LCSS is partly responsible for its better crosstalk performance. The split screen can be multiplexed between two polarization states during the same half cycle of the stereo frame. Therefore, a longer time is available for phosphor decay associated with one state (one section of the split screen) while the opposing state is addressed (the other half of the split screen).

The values for the Tektronix LCSS appear to be more uniform across the surface area of the shutter than for the Z-Screen. Specifically, the lower crosstalk values at the top of the Z-Screen may be due to the inability of the liquid crystal medium to switch completely (to achieve an equilibrium state) within the 0.5 millisecond vertical retrace time.

Normal and off-axis viewing of the Z-Screen show basically the same behavior (see table 2) as the Tektronix LCSS. Some of the luminance values are the same for corresponding pairs of ON and OFF values for normal angle viewing. This behavior is observed at the center and bottom of the Z-Screen. The trend for the Z-Screen crosstalk ratio for off-normal viewing at 45 degrees is more toward lower values than those for the Tektronix unit. For example, top and bottom portions of the Z-Screen have a left eye crosstalk ratio less than two. There is more bleed through off-axis with the Z-Screen, than with the Tektronix LCSS. This may be due to the double layered structure of the Z-Screen. However, there are no significant crosstalk differences with either unit when viewed normal to the screen.

Notice that in tables 1 and 2 the crosstalk ratios are also dependent on the color transmitted through the LCSS system. The crosstalk ratios are largest for red, and lowest for blue, for both on- and off-axis viewing. This is due to the temporal differences between the decay characteristics of the primary color-producing phosphors and the

wavelength dependence of the interference characteristics of the LCSS system.

In the first case, the decay characteristics are shorter in time for red producing phosphors than for the other two primary color producing phosphors. Even though the present CRTs used in workstations are the so-called short persistence phosphor types, there is sufficient residual light output in the decaying portions that can be observed. Bos [3] shows that the green light intensity from a P22 short persistence phosphor in the interval from 1 to 9 ms is about 20 percent of the amount emitted in the interval from 0 to 1 ms. As a result, the "incorrect" eye sees bleed through from the unwanted image about 1/5 as bright as the "correct" image.

In the second case, optical interference properties of the LCSS system, consisting of the LC medium, quarter wave retarders, and polarizing glasses, are wavelength dependent. Ideally, the system is designed to maintain a flat response over all wavelengths produced by the phosphors [3]. However, variations occur and Tektronix claims that perfect quarter wave retarders in the polarizing glasses cannot be obtained. The transmission characteristics of the various colors through the LCSS system will be affected.

We recommend avoiding a black background. The black background will provide too much contrast with respect to the displayed image so that the bleed through, if any, will become noticeable. A gray background reduces the appearance of bleed through. Also, mixed colors with a high green content will show bleed through more than others. This is particularly noticeable with wire mesh graphics imagery than with continuous regions of color imagery. In any case, the bleed through becomes more apparent for off-axis viewing.

2.2.1 Comments on Measurements

Two factors made these crosstalk measurements difficult to assess: one, the vendors do not provide crosstalk data with the LCSS systems; two, we are not yet able to specify a value of crosstalk that would correlate with various levels of user acceptance or quality of stereoscopic image. Of course, values close to unity are not acceptable. The above are areas for further research.

We are also concerned about the precision of the measured values of luminance. The OFF state luminance values are sufficiently low that the

most sensitive scale of the photometer is used. The last digit of the luminance value is of questionable accuracy but significantly affects the calculated crosstalk ratios. When the ratios are calculated, the final values show large variations because of the uncertainty in the last digit.

2.3 ASSESSMENT

The interocular crosstalk ratio is a display parameter that uniquely describes the ability of the LCSS to provide high-quality 3D stereoscopic imagery. A straightforward photometric measurement technique has been devised that measures the crosstalk ratio. The technique has been used to characterize the performance of two commercially available LCSSs, one from Tektronix and the other from StereoGraphics. Both LCSSs produce about the same quality 3D stereoscopic imagery when viewed normal to the CRT screen. Tektronix has a slightly higher crosstalk ratio but the difference is not sharply noticeable. For off-axis angle viewing, however, the StereoGraphics product does not perform as well as the Tektronix product. Qualitatively, the difference is readily observed in the off-axis angle viewing mode. Observers aware of the split screen in the Tektronix LCSS tend to prefer the StereoGraphics screen, regardless of its off-axis performance. However, both shutters provide acceptable stereoscopic imagery for the mission planning and rehearsal simulation described in section 6.

SECTION 3

STEREOSCOPIC PICTURE GENERATION

Computer-generated stereograms of scenes with different degrees of realism are now feasible when the LCCS is driven by a high performance computer graphics workstation, such as the Silicon Graphics (SG) IRIS 4D/70GT. This section is devoted to how to generate "good" stereograms on a 19-inch workstation monitor for consistently comfortable viewing.

3.1 VISUAL DEPTH CUES

Successful application of computer-generated stereograms requires an understanding of the depth cues used by our visual system to interpret the sensory input provided by our eyes. Stereopsis, by itself, is inadequate for effectively communicating depth in a visual scene. It relies synergistically on other depth cues that derive from human physiology, psychology, and experience. The depth cues are generally classified into two categories: psychological cues and physiological cues. The psychological depth cues are:

- a. Image Size. Bigger images are perceived as being closer than smaller ones.
- b. Linear Perspective. Parallel lines converge at an infinitely distant vanishing point
- c. Interposition. Close objects obscure objects that are farther away.
- d. Aerial Perspective. Atmospheric haze provides a depth cue for distant objects.
- e. Light and Shade. Highlights, shadows, and shading, especially from familiar sources such as the sun, provide clues to an object's shape, which in turn reveals something about its depth. In the absence of other cues, for example, we will interpret shading and shadows assuming that light comes from overhead (as it usually does), and we will conclude that objects highlighted on top and shaded on the bottom are convex. With artificial light sources, the attenuation of light over distance provides another depth cue.

- f. Textural Gradient. The textures of close objects are more detailed. Far objects evidence less detailed texture.
- g. Motion Parallax. Near objects move quickly through our visual field. Distant objects, at the same velocity, move more slowly through the scene.
- h. Familiarity. Familiarity is often overlooked as a depth cue, and yet it is so powerful a cue that it will routinely override all others. We are familiar with most objects through everyday experience. The combination of recognition, knowledge of the object's normal size, and comparison to the size of the image, will immediately scale the whole scene according to experience, and contribute to the perception of depth.

Except for familiarity, these depth cues can be derived from the rules of Euclidean geometry.

The physiological depth cues relate directly to the human visual apparatus, and are extremely important for stereoscopic picture generation.

- a. Accommodation. The degree of muscular effort required to focus the lens of the eye may provide a feedback mechanism for distance.
- b. Convergence. In order to fuse the two images, each eye must rotate so that the point of interest is imaged onto corresponding points of each retina. Experiments show that this mechanism does not function as a range-finder [4]. The depth cues provided by this mechanism are strictly relative to other points in the scene.
- c. Retinal Disparity. When the eyes are converged on a particular point of interest, all other homologous points of the scene exhibit a positional difference between the two retinae. (See figure 3.1.) It is this retinal disparity, not the convergence, that contains the depth information for the scene.
- d. Stereopsis. The mind has the ability to combine two images into a single perception of the world. How this happens is not known, but the combined image has a greater sensation of depth, solidity, and texture, than the image from either eye alone.




Figure 3.1. Retinal Disparity of Point B when the Eyes are Converged on Point A is $|d_l - d_r|$

3.2 VIEWING CONSIDERATIONS

In addition to coordinating the proper use of visual depth cues, the programmer must be aware of several constraints that have a direct bearing on the viewability of the stereogram. If these constraints are not considered, it is very likely that the viewers will be uncomfortable, or will not experience stereopsis. Tolin [5] discusses these considerations as "pitfalls" for the display designer; Lipton [6] characterizes them in terms of binocular symmetries.

Parallax is a primary concern in computer-generated stereograms. Parallax is defined as the distance between homologous (corresponding) image points in a stereo-pair, as measured on the display screen. The various types of parallax are described in table 3. Divergent and vertical parallax represent completely unnatural situations, and are likely to be painful for the viewer. The other types of parallax, when mapped to the appropriate eye by the image selection mechanism, correspond to the retinal disparity, which contains the stereoscopic depth information.

Table 3. Parallax	
Horizontal	Horizontal (with respect to viewer) distance between homologous points, measured at the display screen
Zero	lines of sight intersect at screen
Positive	lines of sight intersect behind screen
Divergent	positive parallax greater than interocular
Vertical	Vertical (with respect to viewer) distance between homologous points, measured at the display screen

There is an additional parallax constraint required to assure comfortable viewing conditions. Our eyes have become accustomed to focusing and rotating in a coordinated fashion. Stereoscopic displays, however, require that our eyes remain accommodated (focused) on the display screen, while they converge on a point of interest whose apparent depth may be different than the display screen. Under these conditions, Valyus [7] finds that most people will tolerate a change in convergence angle of up to 1.6 degrees, (figure 3.2). If this tolerance is exceeded,

it is called an *accommodation/convergence conflict*. Valyus' guideline is:

$$|P_{\max}| = 0.03 \cdot D_{\text{view}}$$

where P is the maximum allowable parallax, and D is the viewing distance to the display screen.

The parallax restriction places a limit on the recommended depth range of the stereogram. A simple relation that gives the recommended depth range for a particular viewing distance, based on Valyus' 1.6 degree tolerance, is derived. Referring to figure 3.2, assume for simplicity that the viewer directly faces the display screen. Based on the change in convergence angle of

$$\theta = \alpha - \beta$$

we derive

$$D_{\text{point}} = \frac{T_e}{\frac{T_e}{D_{\text{view}}} - \theta}$$

If we set θ to the maximum change in convergence angle, then for any given viewing distance, we can calculate the minimum ($\theta = -1.6$ degrees) or maximum ($\theta = +1.6$ degrees) recommended depth range. The maximum depth will become infinite when

$$\frac{T_e}{D_{\text{view}}} = \theta$$

By using a maximum change in convergence angle of $\theta = +1.6$ degrees, and an interocular distance of 65mm, we find that a viewing distance of at least 2.3m is required for viewing objects at a distance stereoscopically. If the 2.3m viewing distance is used in the equation for D_{point} with a change in convergence angle of $\theta = -1.6$ degrees, we obtain a minimum recommended object distance of 1.2m. By adjusting the depth range and viewing distance according to this formulation, the parallax will always remain within acceptable limits. The resultant stereoscopic image will be comfortable to view.




Figure 3.2. Accommodation and Convergence. The convergence angle for point A is α , and point B is β . The parallax for point A is 0, and point B is P_B . When viewing the point of interest B, at distance D_{point} , the eye must remain accommodated for the display surface at distance D_{view} .

Another important viewing consideration is maintaining the normal depth relationships between the objects in the stereoscopic scene. This is called *orthoscopy*. Perfect orthoscopy can only be achieved for one viewer at a time, because the perspective projections depend upon the exact viewing position. Except for some entertainment and artistic applications, regard for the viewer requires that we strive for orthoscopy whenever possible.

3.3 CURRENT METHODS FOR STEREOSCOPIC PROJECTION

Stereoscopic projection is generally modeled after stereophotography. Within this class of models, one of two methods is usually employed: crossed or parallel axes projection.

3.3.1 Crossed Axes Projection Method

The crossed axes projection method is illustrated in figure 3.3. In this method the optical axes of the "cameras" cross in space at a point of interest. The image is then projected onto the image plane. This has been the preferred method for stereo-cinematographers.

(A variant description of this method, often found in computer-oriented literature, describes one of the stereo-pair images as being obtained from the other by a rotation through a few degrees of the point of interest. The rotation method is equivalent to the crossed axes method.)

3.3.2 Parallel Axes Projection Method

Figure 3.4 illustrates the parallel axes projection method. This method was not generally used until Baker [8] recommended it in 1987. In this model, the camera axes are parallel in the horizontal viewing plane. The interaxial distance controls the relative amount of parallax in the stereoscopic scene. In addition, the projected images can be shifted horizontally to further manipulate the parallax in order to remain within Valyus' accommodation/convergence guideline. The shifting also affects the apparent depth composition relative to the display screen. Programs using this method sometimes provide the user with several interactive controls: an interaxial separation adjustment; a horizontal image-shift adjustment; and a field of view, or "zoom lens" adjustment.




Figure 3.3. Crossed Axes Projection. The "cameras" are separated by an interaxial distance T_c corresponding to the convergence angle θ . The image planes are at a focal length F_c from the centers of projection. (Note: With actual cameras, the center of projection corresponds to the lens. The lens inverts and focuses the image on a focal plane *behind* the lens. The photographic printing process restores the image to its upright orientation for viewing. For simplicity in a computer model, use the viewing geometry illustrated here.)




Figure 3.4. Parallel Axes Projection. The "cameras" are separated by an interaxial distance T_c . The image planes are at a focal length F_c from the centers of projection. The optical axes are parallel.

3.3.3 Shortcomings of Current Methods

Although both methods can produce stereopsis, in many cases their application in computer-generated CRT displays will lead to an unacceptable result.

The crossed axes technique, for example, will introduce rotational distortions. This distortion produces a "keystone" effect that results in the introduction of vertical parallax in the stereo-pair images. This procedure also produces a warped perception of the scene when viewed stereoscopically because the linear relationship between parallax and depth, expected by our perceptual system, will have been circumvented. The crossed axes techniques has enjoyed some success in the motion picture industry because it is easier to set up the cameras this way, and the distance to the object being photographed is large compared to the interaxial, and the viewing distance is large compared to the interocular. Under these conditions the convergence angle is small, and rotational distortions become insignificant. For typical CRT displays, however, the viewing parameters are quite different, and the distortions often become noticeable.

The parallel axes technique eliminates distortions due to rotation of the scene, but introduces some problems of its own. Figure 3.5 illustrates the situation when viewing these stereograms on a CRT display. The left and right views are time-multiplexed on the display. The action of the LCSS and the polarizing glasses (not shown) maps the appropriate image to each eye. Figure 3.5 shows how the interposition depth cue will be incorrect at the edges of the scene. The area of the CRT that bounds the image is called the *surround*. Some objects apparently behind the display surface on the left side of the scene may overlap the surround. The natural expectation is that the right eye should be able to see around the left edge more than the left eye. The parallel axes projection method, however, produces precisely the opposite effect. No amount of image shifting can correct this defect. There will always be an alignment problem at the edges.

The choice of subject matter affects the suitability of the current techniques. Most computer applications to date have displayed objects floating in space against a black background (graphs, plots, and charts; molecules; crystals; anatomical features; mathematical models of fluid flows; and machine parts). Although these images have real-world analogues, they do not correspond closely to everyday experience. A certain amount of distortion may go undetected, or even be desirable. An object floating in space does not extend to the edge of the surround, and so edge problems will be observed. Some of the most valuable applications of computer-generated stereograms will require a visual simulation of real-world scenes that fill the entire display area.




Figure 3.5. Viewing a Parallel Axes Projection Stereogram. The viewing distance from the display is D_{view} . The interocular is T_e . The shaded volume I is the stereoscopic field of view. Volumes II and III exhibit an interposition depth cue contrary to normal experience. If the focal length or interaxial used in the projection differ from D_{view} , or T_e , the perspective and/or disparity in the stereoscopic image will differ from normal experience, resulting in a distorted perception of depth.

For example, out-the-window views would be required for simulations of aircraft landings, in-flight refueling, handling of hazardous materials using mechanical arms, the space shuttle mechanical arm, spacecraft dockings under manual control, harbor piloting of ships, and remotely piloted vehicles.

3.4 STEREOSCOPIC WINDOW PROJECTION

Camera models are not always appropriate for producing stereograms. Therefore we propose a new model, based on the common experience of looking through a window, called the *stereoscopic window projection*.

3.4.1 Description

We can model the CRT display as a flat rectangular surface that defines a window through which we can view computer-generated object images. The field of view through the window depends on the angle subtended by the display screen from the viewer's position (see figure 3.6). The stereogram is generated by projecting the objects in the scene onto the display surface for each eye. The significant difference from previous methods is that the image plane has the same position and orientation for both eye projections. The center of projection for each view is determined by the positions of the left and right eyes when viewing the scene.

3.4.2 Discussion

This projection contains no rotational distortions, and there is no possibility of vertical parallax. Since the centers of projection are determined by the positions of the eyes relative to the image plane, there can be no divergent parallax. The perspective will be correct, and the interaxial separation will be equal to the interocular. This method will create an orthoscopic stereogram.

The only remaining consideration is the possibility of an accommodation/convergence conflict, that in turn depends on the depth content of the scene and the viewer's distance from the display. If the depth content of the scene or the viewing distance cannot be altered to remain within the limits recommended by application of the equation for D_{point} (section 3.2), there is one acceptable departure from the orthoscopic condition. A scale factor may be applied to the projected image. This operation is called a frame magnification. The magnification factor should be in the semi-closed range $(0,1]$. (With very large screen displays, as in a large movie theater, viewers will actually tolerate a




Figure 3.6. Viewing a Stereoscopic Window Projection Stereogram. The shaded volume I is the stereoscopic field of view. Volumes II and III exhibit an interposition depth cue consistent with experience. Volumes IV and V, due to the finite extent of the display screen, are non-stereoscopic. Since the projection is computed based on the actual viewing distance D_{view} and interocular T_e , the perspective and disparity are consistent with normal experience, resulting in an undistorted perception of depth.

maximum total divergence of 1 degree, which permits magnification factors greater than one. For CRT applications however, divergence should always be avoided.) We can use Valyus' guideline to determine an appropriate magnification factor based on the viewing distance:

$$P_{\max} = M \cdot T_e = 0.03 \cdot D_{\text{view}} .$$

In the stereoscopic window method, the maximum parallax for distant objects will be equal to the interocular. For example, with a viewing distance (D_{view}) of 1.0m and an interocular (T_e) of 65mm, an appropriate magnification factor (M) for a scene containing distant objects would be $0.03 \cdot 1.0 = 0.065 = 0.46$. The magnification factor is acceptable because our visual perception has no absolute sense of image size. As long as the relative disparity remains proportional to the image size, the perception will remain undistorted.

Camera models of stereoscopic viewing obfuscate the geometric considerations surrounding computer-generated stereograms. The projection geometry must be consistent with the viewing geometry. The stereoscopic window model has been implemented as a method for producing stereograms that are comfortable to view and compatible with viewer expectations. Geometric fidelity alone, however, is inadequate for consistently producing successful stereograms. When stereoscopic depth cues are combined with appropriate monocular depth cues, the combined synergism creates a startlingly realistic scene. We use these principles to construct out-the-window terrain scenes for route planning purposes.

SECTION 4

GEOMETRIC MODELING OF TERRAIN

Figure 4.1 is a diagram of the steps involved in transforming digital terrain elevation and digital feature databases into a geometric model suitable for presentation in 3D stereo. Also included is a discussion of techniques for the development and incorporation of generic objects (icons) representing selected point features from the digital feature data into the model, and for representing additional features such as radar intervisibility contours into the model.

4.1 PROCESSING OF DIGITAL TERRAIN ELEVATION DATA

The digital terrain elevation data is processed and the significant terrain features are extracted and merged with the digital feature data. A block diagram (figure 4.2) shows the steps involved in processing the digital terrain elevation data.

4.1.1 Description of Digital Terrain Elevation Data

Digital Terrain Elevation Data (DTED) is a Defense Mapping Agency (DMA) product used throughout the C³ community. It is composed of a matrix of elevation data within 1 degree latitude by 1 degree longitude quadrangles. The data has a precision of 1m and an accuracy of $\pm 30\text{m}$ vertically. The spacing between adjacent samples varies with the latitude, and is approximately 100m between elevation samples. The actual area covered and the number of samples in each quadrangle also varies with latitude (see table 4.1).

Table 4.1 Number of Samples in DTED Quadrangles

<u>Latitude of Quadrangle</u>	<u>Number of Samples</u>
0-50°	1442401
50-70°	721801
70-75°	481601
75-80°	361501
80-90°	241401

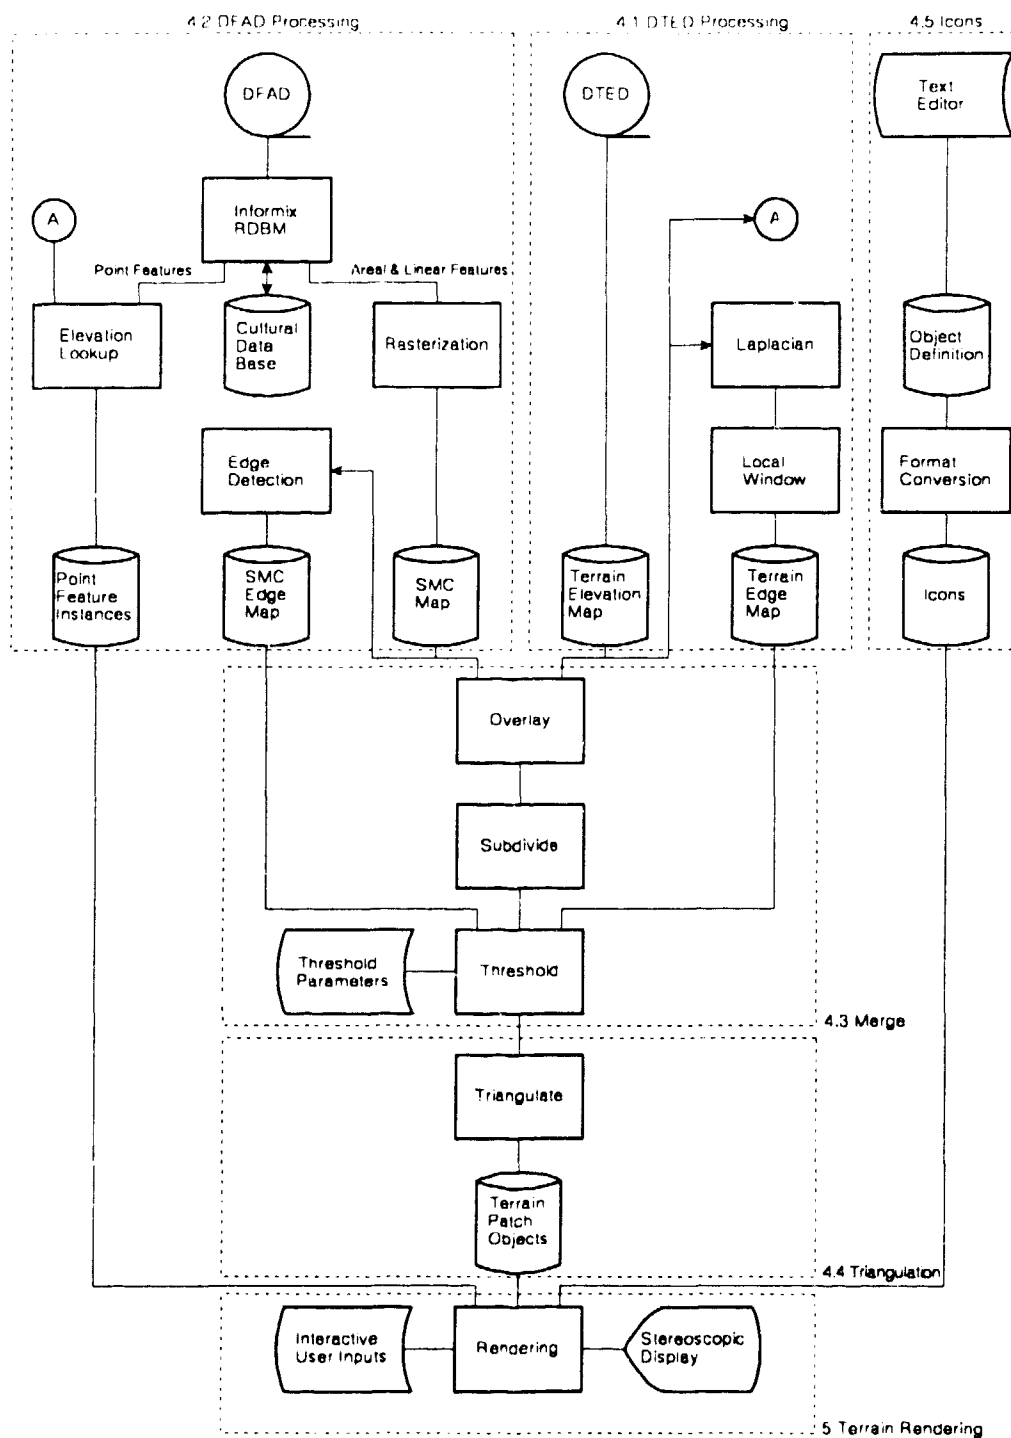


Figure 4.1. Transforming DMA Data: Overview

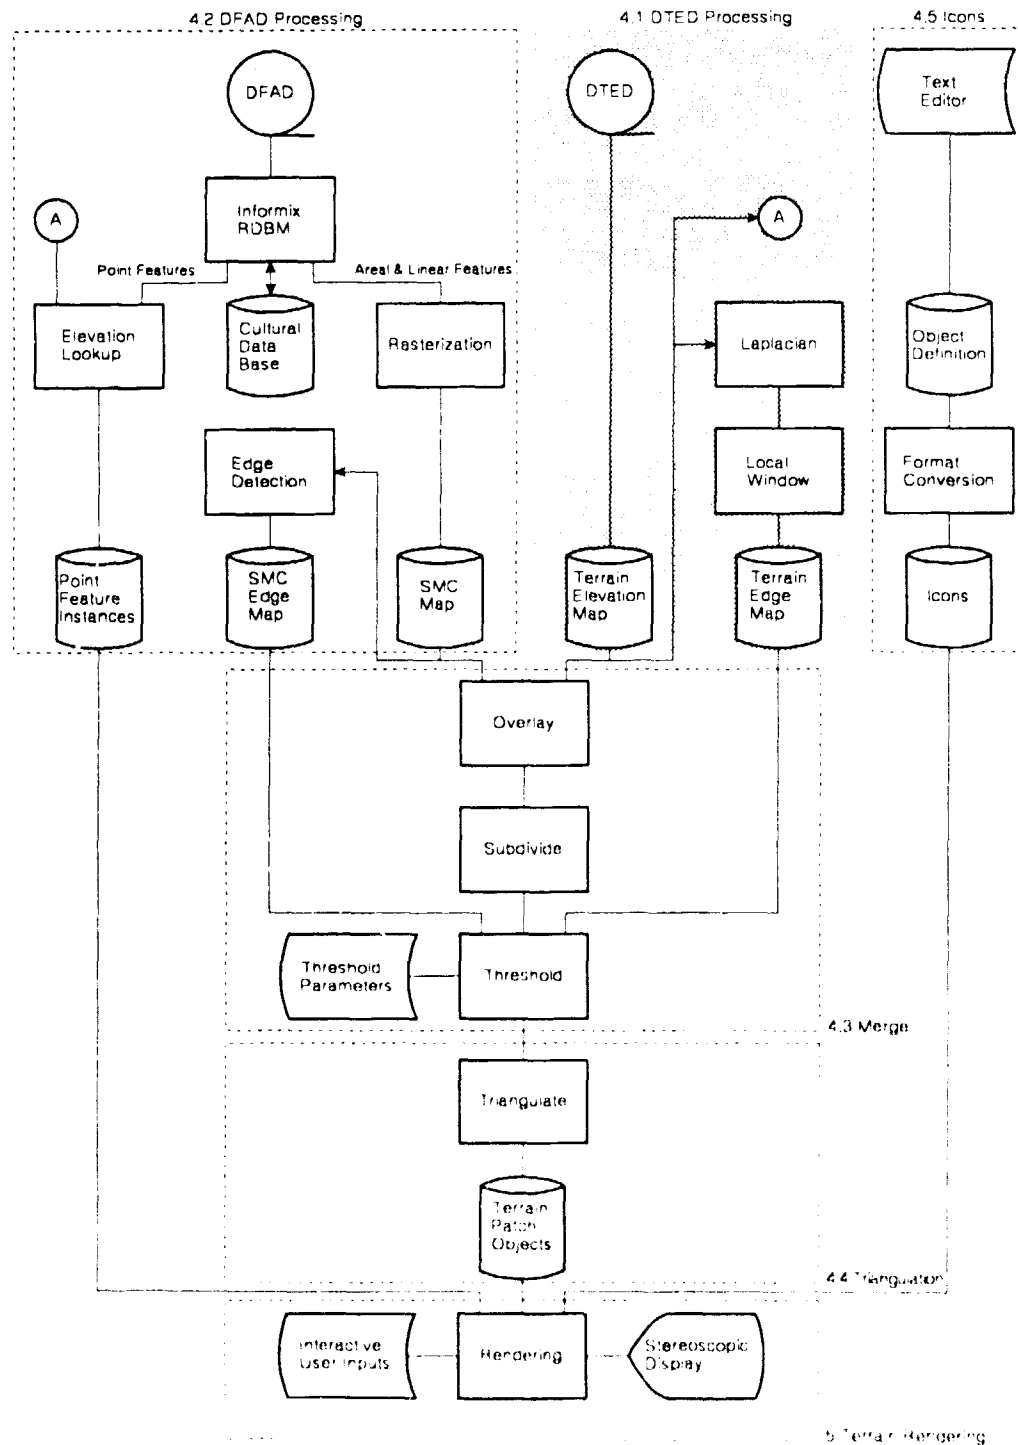


Figure 4.2. Terrain Analysis: Processing of DTED Data

To graphically display only one quadrangle would require a considerable amount of processing. If we used a planar triangular facet as the graphics primitive, a quadrangle in the 50 to 70 degrees latitude range would require 1.44 million triangles. A practical C³I system would require access to several quadrangles simultaneously. Many attempts at generating interactive terrain scenes have been hampered by this large volume of data. (For examples, see references 9-11.)

The simplest way to reduce the amount of data is to choose only every Nth datum from the original matrix. This is quite arbitrary, however, as there might be interesting features between the selected points. Examining the origins of the DTED data, Peucker et al. [12] conclude "most of the DMA grids are highly redundant in terms of apparent average height precision."

4.1.2 Terrain Feature Extraction

Computer-based graphics systems are based on one of two different types of data structure. One is the image data structure, describing pictures in terms of a rectangular array of numbers, called pixels. Each pixel specifies the brightness of a "picture element". This type of data model is appropriate for processing digitized photographs. A primary drawback to this data structure is the large quantity of data that must be maintained. Since the data consists of rectangular arrays, the amount of data increases as the product of its linear dimensions increases.

An alternate data structure is geometric. It does not describe the appearance of the final image, but rather the collection of objects contained in the image. The picture is constructed of geometric primitives, such as line segments and polygons. By operating only on the end-points or vertices of the primitives, a considerable amount of flexibility is achieved. Simple objects may be described with only a few polygons, and complicated objects with proportionally more polygons. The final image will still consist of pixels, but the pixels need not be calculated until the last step. Many stages of manipulations can be performed on a reduced set of numbers. It is practical to interact with the various objects, and consequently to modify the picture contents and viewpoint quickly.

The DTED data essentially follows the image data model (a rectangular array of samples). A more efficient representation, in terms of quantity of data and time required to manipulate it, would be a geometric model. The first step in building this geometric model will be to define what kinds of features are important for modeling terrain. We will borrow some

techniques from image processing technology to extract these features and input them into geometric model algorithms.

The important features for terrain data are:

- a. peaks
- b. ridges
- c. passes
- d. valleys
- e. depressions
- f. changes in slope

These features characterize any terrain surface. We discover which points from the elevation sample matrix lie along these kinds of features and eliminate the remaining points as redundant. The remaining points are derived by a linear interpolation between the surrounding features. The terrain could then be drawn as planar facets, whose edges are determined by the feature points. A common component of all these features is that they represent a curvature in the surface.

4.1.3 Laplacian Operator

The Laplacian operator is a second-order derivative operator and can be used to detect curvature in a surface. The Laplacian is one type of edge-detector commonly used in image processing applications [13]. It is a scalar-valued operator, whose value is zero for functions with constant slope, and non-zero for functions with changing slope. This operator is ideal as a feature- or edge-detector for terrain elevation data.

4.1.4 Adaptive Filtering

In digital image processing, the edge detection function can be combined with the original image to form an enhanced image or, a thresholding function can be used to produce an edge map. The thresholding operation compares the absolute value of the function to a threshold value. If the functional value at a certain point is greater than or equal to the threshold, that value is retained in the edge map, otherwise the point is rejected from the edge map.

Using the Laplacian operator, the edge map will correspond to the high curvature terrain features. The selection of the threshold corresponds to the amount of curvature that is significant to the application. To

determine the threshold value appropriate for a particular application, consider the smallest curvature feature to be detected. However, selecting a constant threshold does have the potential of producing large inaccuracies in regions characterized by small curvatures. In order to prevent these omissions, an adaptive filter has been used prior to the thresholding stage. This adaptive filter is labeled "Local Window Filter" in figure 4.2.

This filter uses a 7 x 7 sliding window over the Laplacian filtered data set. The value of the central component of the window is ranked with respect to its surrounding values. This produces an enhanced terrain edge map based on the rank value for each point, instead of the actual Laplacian value. Since the rank value is determined locally, this filter adapts to the curvature values found in each region. The adaptive filter technique provides a more reasonable representation of the original DTED data set than the constant threshold approach, and with approximately the same number of sampled data points.

4.2 PROCESSING OF DIGITAL FEATURE ANALYSIS DATA

Digital Feature Analysis Data (DFAD), also from DMA, complements the DTED data by supplying the cultural features. This information enhances the terrain model by demarcating areas of forest, rocks, water, and urban development, as well as specifying the size and location of significant point features such as buildings, steeples, and smokestacks. By adding DFAD data to DTED data, a terrain model with a high level of realism can be achieved. The steps in DFAD processing are highlighted in figure 4.3.

4.2.1 Description of Digital Feature Analysis Data

The DFAD database is a collection of cultural features derived from various planimetric, photographic, and intelligence sources. For certain areas of interest, DFAD data may be provided at two data density levels. Level 1, is a generalized description intended to cover large areas of the earth's surface, and level 2, is a detailed description intended to cover small areas of interest. DFAD level 1 data were employed for this effort. (Extension of these techniques to DFAD level 2 data would be straightforward.)

The DMA DFAD database consists of records. Each record describes a single feature. There are three different types of records used to describe different classes of features:

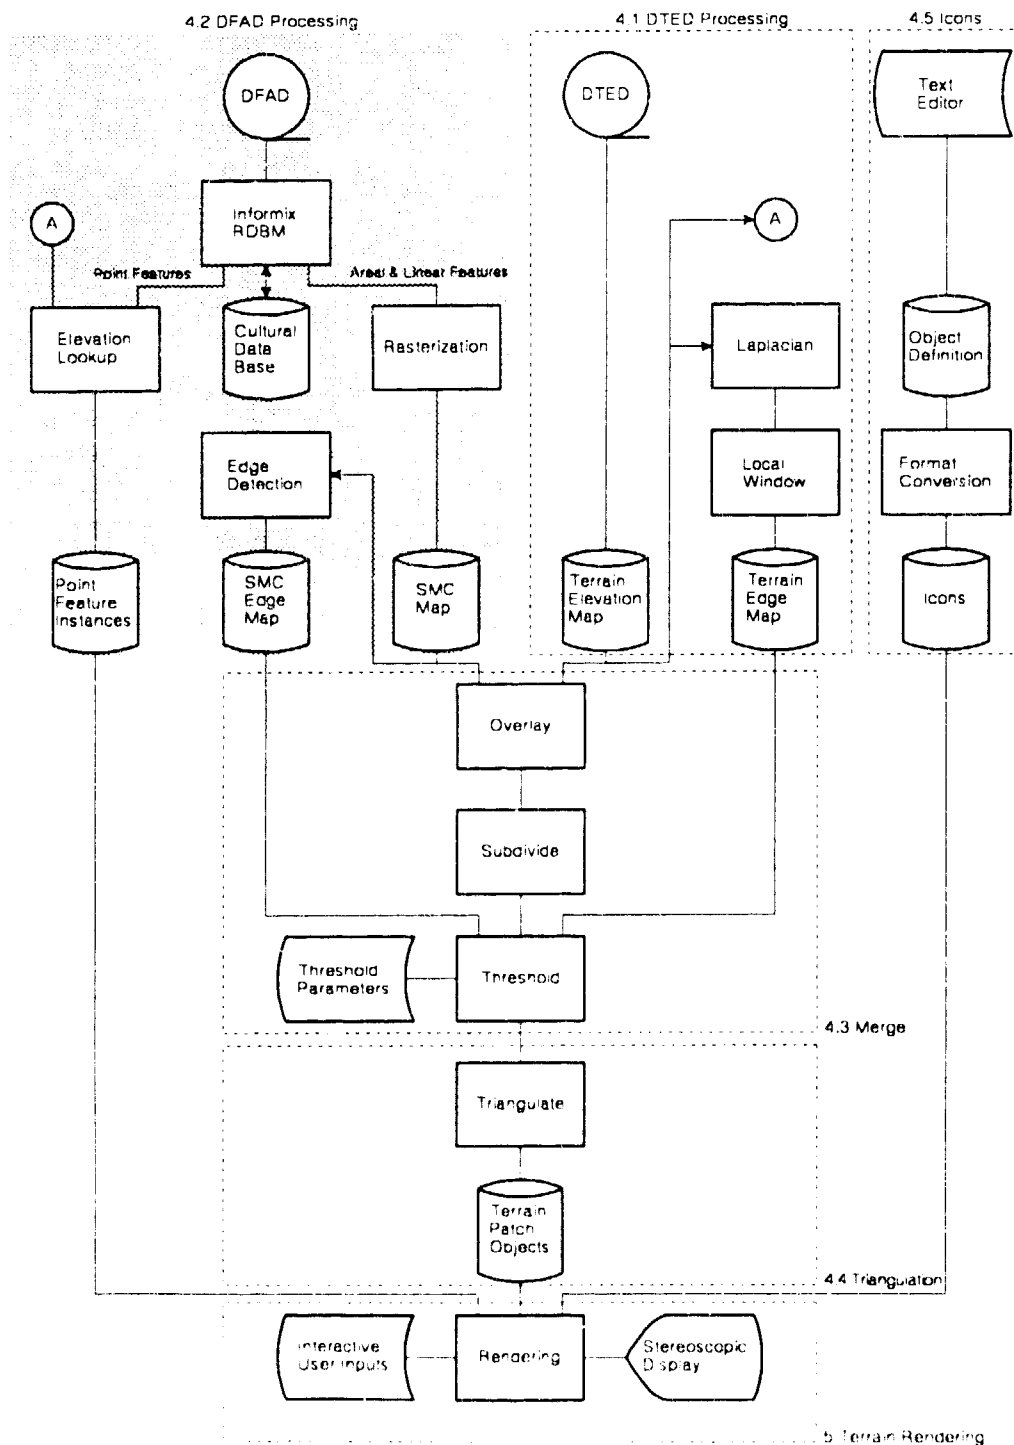


Figure 4.3. Feature Analysis: Processing of DFAD Data

- a. Point. Relatively small feature whose location can be identified by a single latitude/longitude point (e.g., building).
- b. Linear. Long-thin feature whose location and shape can be described by a sequence of connected latitude/longitude points (e.g., railroad track, road, small river).
- c. Areal. Large feature whose extent can be described by the area enclosed by a sequence of latitude/longitude points (e.g., forest area, industrial complex, lake).

In addition to latitude/longitude points, several other pieces of information may be specified for each feature including:

- a. Surface material code (SMC). A number from 1 to 14 which identifies the material from which the feature is composed (e.g., water, metal, stone, concrete, asphalt).
- b. Predominant Height. Height of the feature above ground level.
- c. Feature Identification Code (FID). One of 336 different codes that identify the type of feature (e.g., deciduous trees, silo, school, causeway, smokestack, church steeple).
- d. Length. Length of point features.
- e. Width. Width of point and linear features.
- f. Orientation. Orientation (compass heading) of point features.
- g. Percent Roof/Tree Coverage. Amount of areal feature covered by roofs or trees.

4.2.2 Features Database Management

In order to exercise control over which features from the DFAD database are incorporated into the terrain model, the DFAD data is converted into a relational database under the Informix database management system (DBMS). This procedure involves several steps to convert the DFAD tapes to an ASCII format. Once converted, the DFAD data is further processed and split into two linked data sets which roughly correspond to "feature descriptions" and "feature coordinates" of the DFAD features. These data sets are then loaded into the Informix software as two linked relational databases.

Subsets of the DFAD data are extracted using Informix Relational Structured Query Language (RSQL). The Informix software acts as a filter for selecting features from the DFAD data. We developed utilities using RSQL which allow DFAD features to be "filtered" based on a combination of geographic position, extent, FID, SMC, or any of the other DFAD feature characteristics. This provides great flexibility in controlling the level of realism in the terrain model. (Note that while greater realism is desirable, it reduces the frame update rate during a dynamic rendering of the data.) The output of this filtering process is a formatted ASCII file containing the selected data. We selected all features within the area of interest for our mission planning and rehearsal demonstration.

4.2.3 SMC Map and Icon Instancing File Construction

After the DFAD feature data has been filtered by the Informix RSQL routines, it must be processed into a form that will allow it to be merged with the DTED terrain elevation data. An off-line utility processes the filtered DFAD features and produces two output files. At this stage, the three different types of DFAD features (point, linear, and areal) are treated differently.

The areal and linear features are rasterized into a single SMC image file which can be "draped" over the DTED terrain. The SMC image is a 2D surface material map (figure 4.4). The combination of the SMC map with DTED data produces a complete topographical model of an area including altitude, vegetation, and urban areas.

Point features are processed and recorded as individual occurrences, or instances, in an icon instancing file. This data produces an itemized listing of selected small features, such as houses, steeples, and smokestacks, that are combined with the topographical model to produce a realistic landscape model.

Selection of point features for instancing is based on the FID code associated with each feature that defines the type of point feature (e.g., smokestack, steeple, building). Another means of designating the data to be rendered in the final scene is by selecting only certain FID codes for inclusion in the icon instancing file. The latitude/longitude position, width, height, orientation, and surface material are recorded in an instancing file for the point features selected. The instancing file is further processed to include altitude above mean sea level for each feature in order to specify the location of the features in 3D space. This information is obtained from the terrain elevation map.

Figure 4.4. Example of 2D Surface Material Map



Features with the 11 most predominant FID codes were selected for inclusion in the mission planning and rehearsal simulation instancing file. These 11 FID codes accounted for 95 percent of all point features available in our database.

4.2.4 SMC Map File and Edge Map File Creation

The SMC map is very similar to a satellite photograph. It consists of an array of points, or pixels, that covers a user-selectable range of latitude/longitude values. Each point in the array contains a value that identifies the surface material (trees, water, rocks) at that latitude/longitude point.

The SMC map is created by using a graphics library to rasterize the features into a frame buffer and then dumping the contents of the frame buffer. The size of the SMC map will directly affect the minimum resolution of the resultant terrain scene for a given range of latitude/longitude values. To provide flexibility in map size as well as allow portability among host machines, the MIT X-Window system graphics interface performs this rasterization. The areal and linear features are converted into 2D polygons and polylines, respectively. For linear features, the DFAD "width" attribute is ignored. The linear features are rasterized with a width of one SMC image pixel.

In a fashion similar to the processing of the DTED data, an edge detection filter is applied to the SMC image data to extract those points where the surface material changes. This process selects the important points from the SMC image and creates an SMC edge map.

The parameters used in generating SMC maps for the demonstration produced a 601-pixel x 1201-pixel SMC map for each quadrangle, corresponding to the associated DTED data.

4.3 MERGE OF DTED AND DFAD DATA

In order to merge the DTED and DFAD data, they must be brought to a common ground. This is accomplished by rasterizing the DFAD data into an SMC map. Furthermore, the corresponding DTED and DFAD edge maps provide a way to select terrain surface features and material boundaries. It is now possible to merge the two types of data into a coherent terrain model. Figure 4.5 shows the data flow and processing steps involved.

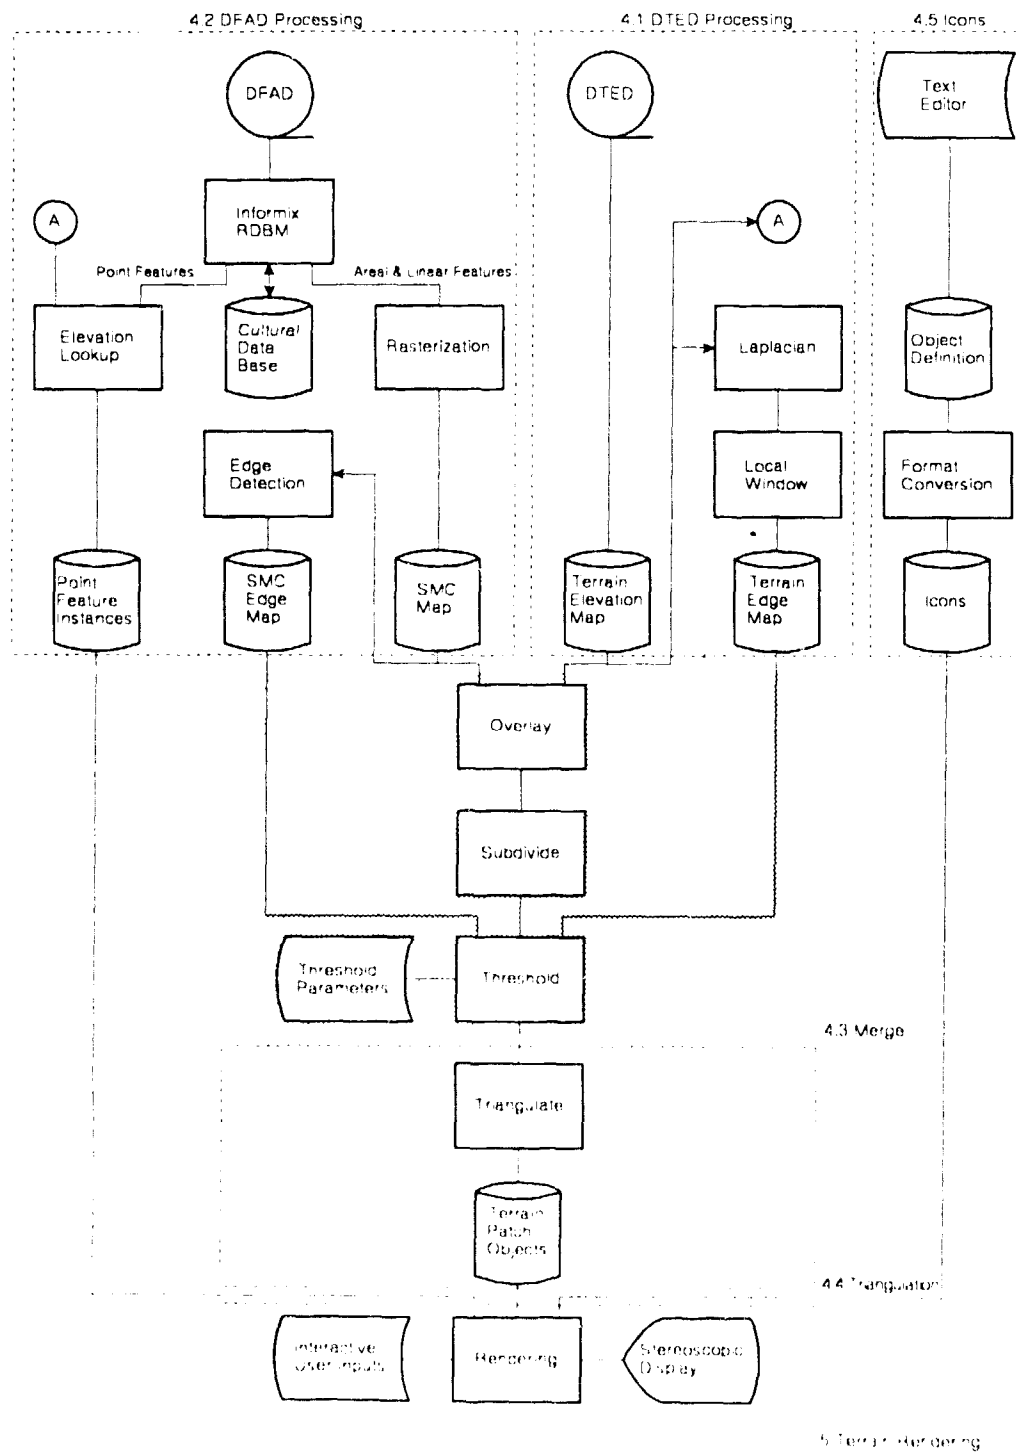


Figure 4.5. Terrain Modeling: Merge of DTED and DFAD Data

4.3.1 Generating the Combined Elevation/SMC Map

Once the DFAD data has been processed into an SMC map, it can be combined with the original DTED data. These two grids of points are registered on top of each other, and a combined grid is generated. The DTED data, SMC map, and combined grids, may have different resolutions. The software performs bilinear interpolation as necessary, on the DTED and nearest neighbor selection SMC data, to produce the combined elevation/SMC map.

4.3.2 Subdivision

The combined grid, SMC edge map, and DTED edge map are then integrated into a geometric model. Instead of combining these three data sets into one huge geometric model, they are combined into small sections. Each of these small pieces, or patches, is three arc minutes of latitude by four arc minutes of longitude. The combined elevation/SMC map quadrangle is subdivided into 300 such patches.

4.3.3 Thresholding

Significant vertices will be extracted from the combined grid for each patch. This process is accomplished by performing a thresholding function on both the terrain and SMC edge maps. This operation compares the absolute value of each edge map to a threshold value. If one of the edge map values at a certain point is greater than or equal to its threshold, that point is considered significant and is retained, otherwise, the point is rejected. If a point is selected as important in either the SMC edge map or the DTED edge map, then that point will be included in the geometric patch model.

The thresholding step is repeated three times for three different thresholds values. This produces three different sets of significant vertices which will model the terrain patch at three different levels of detail. These sets are hierarchical in that the vertices of the lowest-detail set are a subset of the middle-detail set, that is in turn a subset of the highest-detail set. During the fly-through, one of these three models is selected by the rendering software to render a particular patch, depending on the distance of the patch from the current viewpoint.

4.4 TRIANGULATION

Up to this point, the various processing stages have selected a set of points that represents the terrain features at various levels of detail. The next step is to transform this collection of points into a set of polygons with vertices at these points. These polygons will model the terrain and will be colored and shaded appropriately, according to the surface material type and lighting model. These processing steps are shown in figure 4.6.

4.4.1 Triangulated Irregular Network Model

The rectangular array data structure is inefficient in terms of quantity of data and processing time required to transform it. Peucker et al. [12] recommend the Triangulated Irregular Network (TIN) data structure for modeling terrain. We have adopted this structure. The TIN model consists of triangular facets that adapt in size and shape to the curvature of the terrain. For example, rolling plains can be represented with a few large triangles, whereas jagged mountainous regions are described by many small triangles. The TIN is stored by noting the location of all the vertices of the triangles. A list of pointers to adjacent vertices is maintained for each vertex. Any three mutually adjacent vertices form a triangle. The number of triangles that meet at any particular vertex, and the size and shape of each triangle, is variable.

If the TIN is digitized directly from the terrain elevation data source, such as stereo-pair photographic analysis or contour maps, it is not only more compact, but also a more accurate representation than a grid model. A TIN can also be derived from a grid structure, such as the DTED grids. As noted in the previous section, the Laplacian operator can be used to detect terrain features in a grid model, but it does not provide any information on how to connect the features to form a TIN.

4.4.2 Delaunay Triangulation

The Delaunay triangulation is a method for constructing a triangulation of arbitrary points in the plane. The criterion for constructing the triangulation is a "local optimization" test that states that given any three vertices of a triangle, no other vertex in the triangulation will fall within a circle defined by the three vertices. Various algorithms for constructing Delaunay triangulations have been proposed. We have selected an algorithm developed by Lee and Schachter [14] that uses a divide-and-conquer approach. The processing time required by this algorithm is $O(N \cdot \log N)$, which is asymptotically optimal for this problem [15].

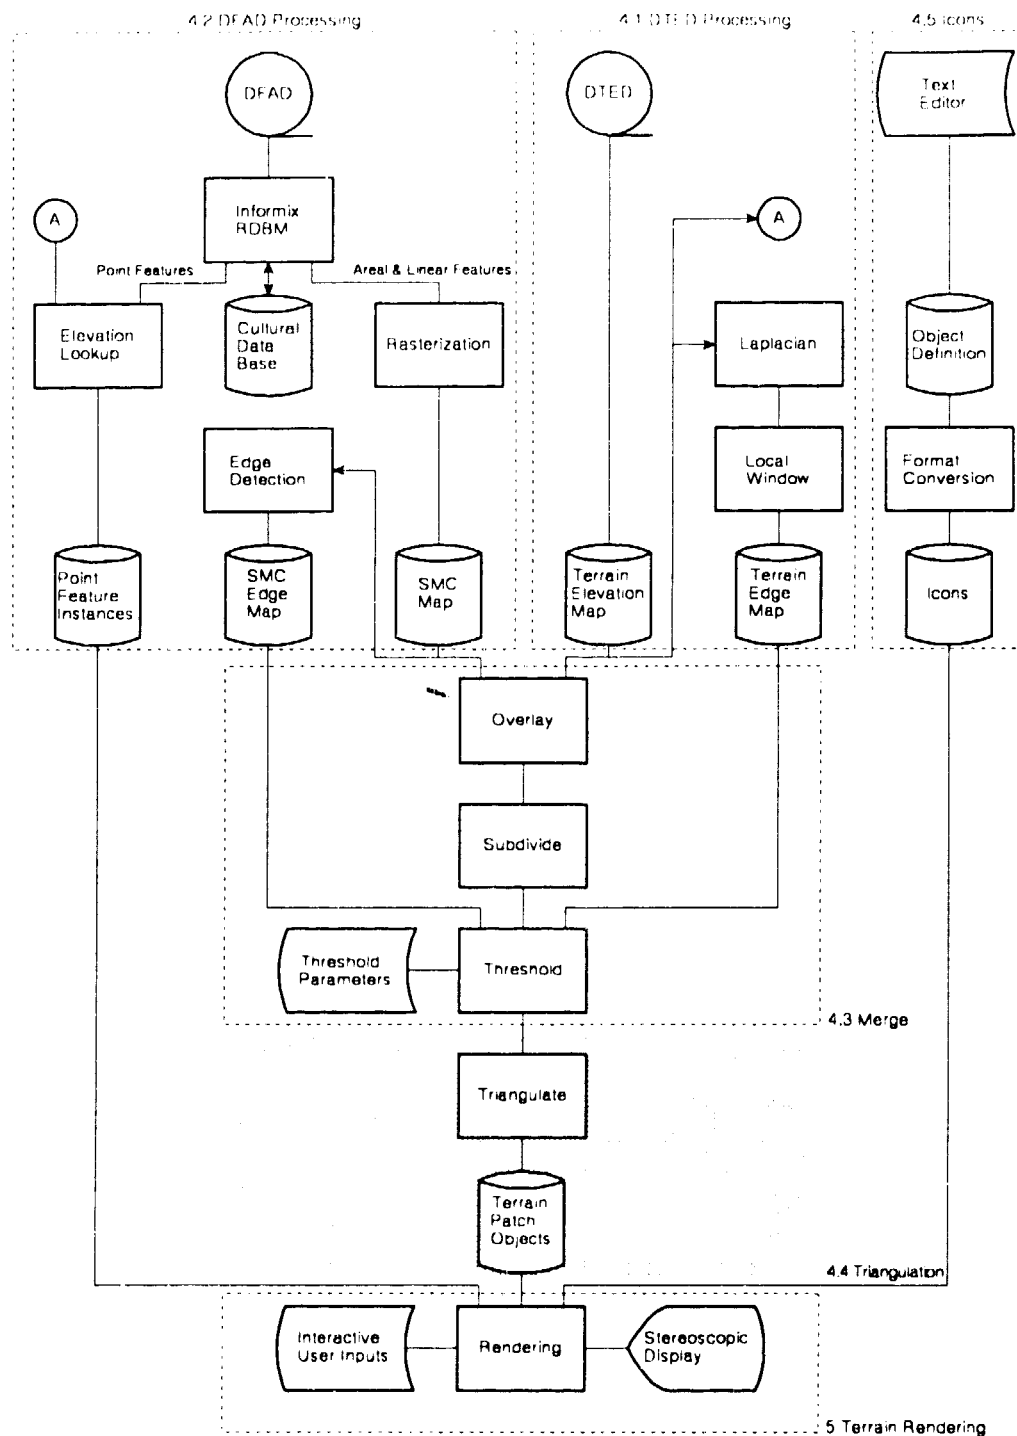


Figure 4.6. Triangulation of Merged Data

4.4.3 Triangular Meshes

Once the TIN has been generated, it can be used to develop drawing routines for the terrain surface. Although drawing may proceed directly from the TIN as individual triangles, preprocessing the TIN into a display list structure of triangular meshes allows faster rendering. This sacrifices some of the compactness of the TIN, but it increases the performance of the rendering process considerably.


Another reason for choosing the TIN model is that our graphics workstation is especially efficient when drawing a batch of connected triangles, called a "triangular mesh". The efficiency arises because adjacent triangles in a triangular mesh share two vertices. Rather than having to transform each of these vertices two or more times (once for each triangle that adjoins the vertex), the triangular mesh primitive allows the vertices to be transformed only once per mesh [16]. The display list generation algorithm recursively descends the TIN data structure, attempting to build the longest triangular mesh it can. It saves each mesh for later use during the rendering process.

In our terrain model, we selected a preprocessing procedure to reduce the total number of polygons required to accurately represent a terrain surface. The accuracy is maintained by employing an adaptive threshold technique for the Laplacian operator. The TIN model reduces the number of polygons required. The triangular mesh graphics primitive reduces the number of vertex transformations required.

Figure 4.7 demonstrates the usefulness of the adaptive threshold technique for terrain visualization. Figure 4.7a shows a 2D plan view of a portion of the original DTED grid utilizing every point. The bright shaded regions indicate larger values of elevation than the darker regions. Figure 4.7b shows a 2D plan view of the same region obtained with a constant threshold. Figure 4.7c shows the adaptive threshold. Comparing figures 4.7b and 4.7c with figure 4.7a, the constant threshold technique does not accurately model the flat, low-lying regions. The adaptive threshold method, however, models not only the hills, but also the flats.

4.5 ICONS

We have described processing steps involved in producing a terrain model including both elevation data, and surface material data. This section will focus on techniques used to place 3D objects into the terrain model. There are two reasons for doing this. The first is to add realism



4.7a

4.7b

4.7c

Figure 4.7. Comparison of Threshold Techniques for DTED Model

to the landscape by providing a means of visualizing selected man-made point features from DFAD. The second is to allow the observer to visualize intangible objects. For example, our mission planning and rehearsal simulation presents radar intervisibility contours.

A flexible "object definition" data structure allows 3D objects to be defined using several different types of graphics primitives and shading algorithms. In addition, each primitive of the object can have individualized surface material attributes. Several format conversion utilities convert these generic object definitions between text files and efficient binary data structures. These steps are shown in figure 4.8.

4.5.1 Representation of DFAD Point Features Using Icons

Although DFAD defines 336 different FID codes that can be used to identify different types of point features, an analysis of our DFAD data showed that the 11 most predominant FID codes account for 95 percent of the point features. By concentrating on these 11 FID codes, the number of point features in the terrain model is maximized while effort involved in developing the 3D icon definitions is minimized.

To further reduce the level of effort involved in developing icon definitions, a single icon can be used to represent several similar FID codes. For example, one generalized "powerline pylon" icon can be used to designate the following DFAD FID codes:

- Powerline Pylon, Type "A",
- Powerline Pylon, Type "H",
- Powerline Pylon, Type "I",
- Communication Tower, General.

Point features with the 11 predominant FID codes can be effectively represented by six icons (figure 4.9). Note that the size and appearance of each particular instance of the icon in the landscape model is controlled by the associated length, width, height, orientation, and surface material from the icon instancing file. This allows the icons shown in figure 4.9 to be scaled, rotated, and "painted", as appropriate, for each DFAD point feature.

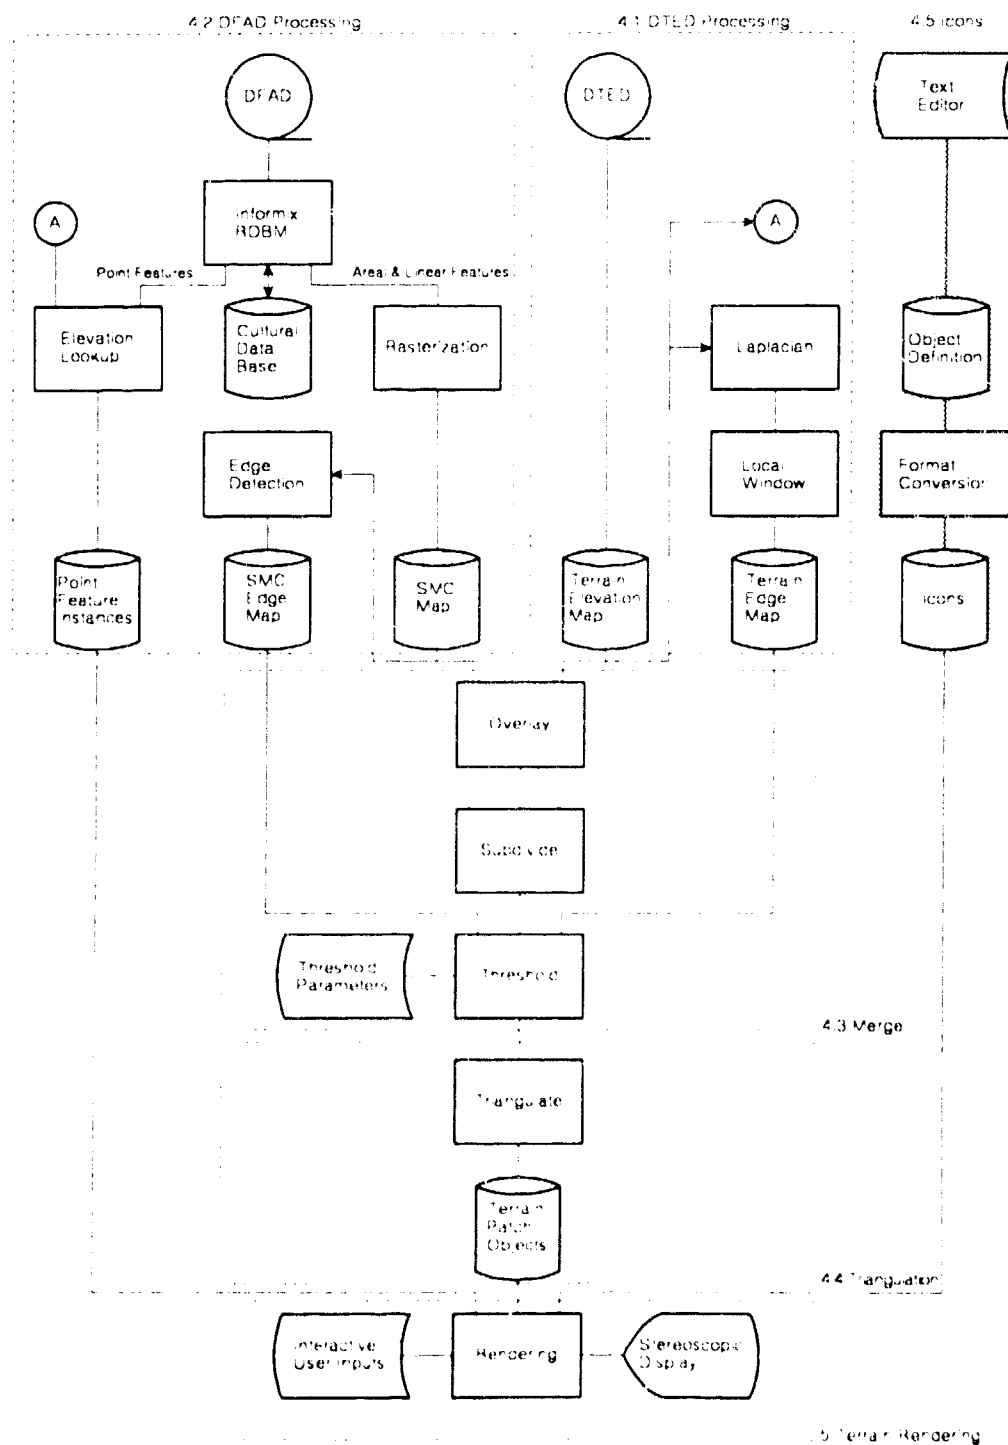



Figure 4.8. Icons




Pylon



House



Steeple



Smokestack



Cube

Gable

Figure 4.9. Icons Developed for Mission Planning

The icons were defined in a manner similar to the terrain patches described above so that they could be rendered at several levels of detail, depending on how far away the icon is from the viewpoint. Figure 4.10 shows these different levels of detail for the powerline pylon icon. The pylon is shown in full detail. As the object gets further away and its size on the screen drops to several pixels, a less detailed rendition is used to reduce the number of graphics primitives needed to render the object.


4.5.2 Intervisibility Contours

Radar intervisibility contours were added to the terrain model to provide assistance in mission planning. The purpose of these contours is to identify those areas in the landscape that are within radar coverage of a particular threat, such as a Surface-to-Air Missile (SAM) site. Since the coverage space of a typical radar is a 3D volume, the intervisibility contours are portrayed by using a wire mesh that defines the volume (figure 4.11).


The shape of the intervisibility contours has been extracted from the Computer Aided Mission Route Planning Aid (CAMRA) software package (previously developed by MITRE under project 411B, Airborne Battlefield Command and Control Center (ABCCC)), and which runs on a MicroVAX II/GPX workstation. The CAMRA software uses DTED data and terrain masking algorithms to define the 3D volume coverage of an idealized radar placed at a particular location in the terrain. The radar model used by CAMRA assumes a beam of fixed range with unlimited azimuth and pitch control. Thus, on a perfectly flat terrain the 3D coverage volume would be a hemisphere. On actual terrain, however, some of the hemisphere is masked due to terrain interactions.

4.6 IMPLEMENTATION

A Silicon Graphics 4D/70 GT workstation was used for all of the processing in section 4 except for the SMC map rasterization. A Stellar GS1000 was employed for the rasterization due to the availability of the X-window system. All of the software is programmed in C. The time required to process the DMA formatted data to a complete terrain model is about 8 hours for a one degree by one degree quadrangle.



Up Close (High Resolution)



Far Away (Medium Resolution)



Very Far Away (Low Resolution)

Figure 4.10. Level of Detail



Figure 4.11. Intervisibility Contour

SECTION 5

TERRAIN SCENE GENERATION

The final processing step, as highlighted in figure 5.1, is to convert our geometric terrain model into a terrain scene display. This process integrates the terrain patch objects, point feature instances, and generic object descriptions, into a complete terrain scene. The results are illustrated in figure 5.2. We will discuss four important considerations with respect to the terrain scene display:

- a. Rendering. Selection of appropriate computer graphics techniques to draw the terrain scene.
- b. Complexity Management. Methods used to organize various scene data.
- c. Performance. Results of measurements based on scenes generated.
- d. Recommendations. Directions for further development of this display technology.

5.1 RENDERING

Reference 1 provides an overview of rendering techniques applied to terrain data. We eliminated inappropriate techniques such as wire-frame and pseudo-color to emphasize realism for the demonstration scenes. An appropriate lighting model is one of the best ways to achieve realism. Our terrain scene uses a Phong lighting model to simulate sunlight. The time of day may be selected by the user.

In conjunction with the lighting model, a shading technique must be chosen. We use one of the following techniques, depending on the type of object being rendered:

- a. Flat shading. For "faceted" objects, such as buildings, where edges and corners must be well defined.
- b. Smooth (Gouraud) shading. To achieve a smooth, curved appearance. For example, the cylindrically shaped smokestack.

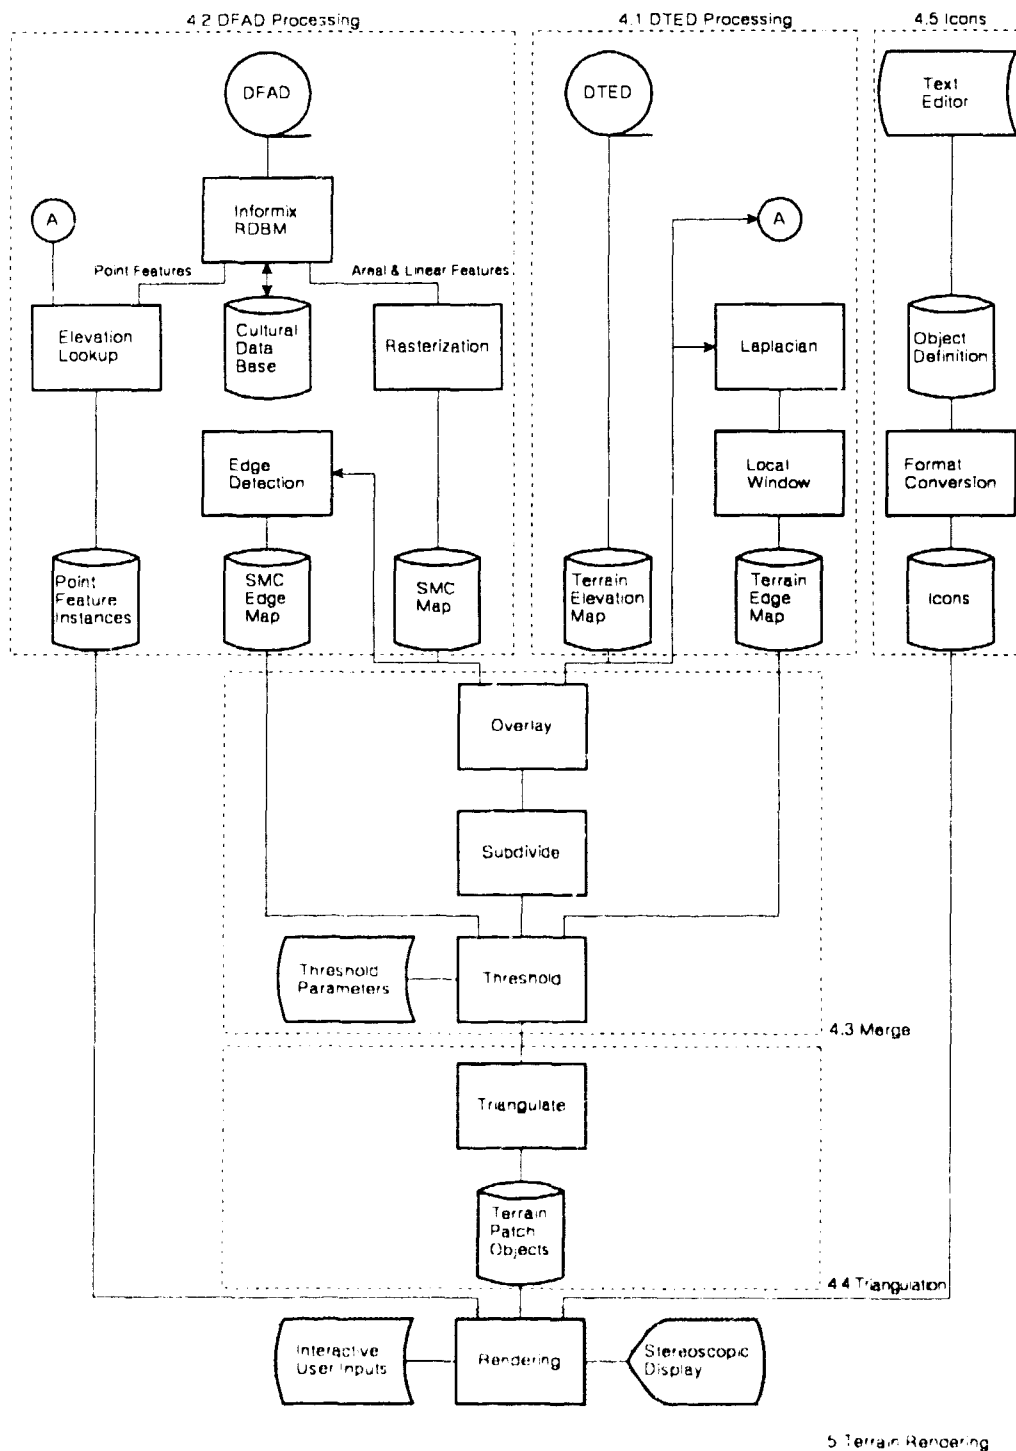


Figure 5.1. Interactive Terrain Rendering



Figure 5.2. Out-The-Window View of Terrain Scene

- c. Blended shading. A variant of smooth shading in which the polygon corners may vary in intensity and color, simulating a transition between different materials. This is the technique used for the terrain surface. The blended technique was judged to produce a better effect for natural objects than smooth shading, because it suggests gradations in texture often found in nature.

The colors were selected for each of the DFAD material codes by performing color measurements on actual materials corresponding to the SMC categories. By calibrating our CRT monitors to known color standards, we are able to produce the measured hues on the display screen.

5.2 COMPLEXITY MANAGEMENT

Terrain scenes are much more complex than most objects considered for computer-generated images. Like most natural objects, a terrain surface cannot be represented by a few simple geometric shapes. In the previous sections we described techniques used to reduce the large number of sample points required to model a terrain surface by selecting an appropriate filter and constructing a network of irregularly shaped triangular facets. These procedures are performed off-line to generate a display list suitable for interactive rendering. Even with this technique however, the amount of data required for a realistically-sized area of interest can overwhelm a workstation's processing capabilities. Consequently, we have incorporated a number of "complexity management" strategies in an effort to maintain a high-quality image and interactive graphics response.

Our complexity management scheme uses the following strategies:

- a. Display list structure. Both the terrain surface and "cultural" structures on the terrain, are handled with the same general purpose display list structure.
- b. Object decomposition. The scene database is divided into many pieces, called "objects." An object is used to describe graphically how to draw each of the pieces. An object may be a small patch of the terrain surface, or a building or other structure. For organizational purposes, each object may be composed of several "parts." Each part, in turn, may be composed of several "primitives," which correspond directly to the graphics primitives available on the workstation: points, polylines, polygons, and triangular meshes.

- c. Generic objects. Generic objects may be constructed and instantiated with the correct size, shape, location, orientation, and surface material, as specified by the DFAD database. (See section 4.)
- d. Field of view culling. All the objects in the terrain scene are tested against the viewer's field of view and the distance to the horizon (see figure 5.3). Objects lying outside the field of view are not considered for further graphics processing. Relegating this decision to the graphics processor's clipping algorithms incurred a significant overhead associated with processing graphic primitives that ultimately had no effect on the picture.
- e. Level of detail. Objects may be defined at several levels of detail to reduce the overall scene-drawing complexity and time. For example, a power transmission pylon on the distant horizon may be adequately represented by a single vertical line segment. However, if the pylon is in the foreground of the picture, its detailed structure must be shown. Each object, whether a terrain patch or building, has several "renditions" associated with a predetermined level of detail metric. During the scene picture generation, a particular metric value is calculated for each object instance based on the current distance from the viewer. This value then determines which rendition is chosen, by comparing the preset generic value. Since there are generally fewer detailed objects in the background, and relatively few highly detailed objects in the foreground, this method results in a significant savings in drawing time.
- f. Memory Management. The display list for the entire area of interest is too large to fit into the workstation's main memory. In a virtual memory operating system, such as the UNIX system used by most workstations, exceeding the memory limits can lead to a condition called "thrashing," in which performance is drastically reduced while the operating system is attempting to swap in-memory data with on-disk data. This problem can be alleviated by releasing the memory taken up by display list data that is no longer in the field of view. This also means that as new regions enter the field of view, the display list must be read-in from disk. These "garbage collection" operations exact a certain performance cost, but it is usually far better to perform these operations explicitly, rather than allowing the operating system to perform them implicitly, because the operating system may blindly swap out data that is still required for picture generation, further compounding the problem.




Figure 5.3. Field Of View Culling

- g. Shared calculations for stereoscopic views. Stereoscopic imagery requires that each view be rendered from a different perspective. We have tried wherever possible to eliminate duplicate calculations. This applies especially to the field of view culling.

Performance profile analysis of actual interactive sessions show that these complexity management techniques are effective. With the current hardware configuration, the processing load is approximately balanced between the CPU and graphics processor unit. The CPU time is spent primarily in graphic library and drawing routines.

Quantitatively, the performance of our demonstration agrees with our previous benchmarks of the IRIS 4D/70GT. The low-level benchmarks show that the IRIS is capable of a maximum throughput of 130,000 facets/second, but only with specialized, highly-tuned programs. A peak rate achievable by practical applications was measured in the range of 38,000 to 62,000 facets/second. In the current demonstration, the scene complexity varies from frame to frame, with an average value of about 25,000 facets. We are able to render these scenes stereoscopically at the rate of one to 2/3 per second (except when memory management operations intervene). This yields a throughput of about 33,000 to 50,000 facets per second.

An animation rate of about one-frame per second is adequate for fly-through sequences with limited "quantized" interactive controls: start, stop, turn left, turn right, reset. This rate however, is inadequate for true "hands-on" interaction because eye-hand coordination requires a minimum rate of five-frames per second. With the current capability, we recommend preset routes through the terrain database, with limited interactive controls. In the near future we expect that workstations with maximum quoted throughputs of about 500,000 facets per second will become available. This rate represents about a five-times performance increase over current models. Continuous "hands-on" control will be feasible with this application, using mouse or joystick devices. If a ten-frame per second rate can be obtained, perhaps through additional improvements in modeling and complexity management schemes, then quite realistic flight simulations are possible. We have been experimenting with a six-degrees-of-freedom input control device called Spaceball, which would be valuable at these enhanced frame rates.

Qualitatively, the scene rendition is quite good at 25,000 facets. An example of such a scene photographed from the CRT display is shown in figure 5.2. Given the fixed peak throughput of the workstation's graphics processor, and assuming that other elements of the program are well balanced, the amount of detail, measured by the number of facets in a

scene, may be exchanged for animation frame rate. Current special-purpose flight simulators actually process only a few thousand facets per scene, but compensate for this limitation with special 3D texture-mapping hardware. Graphics workstations currently available lack this capability, and must compensate by using more facets to obtain adequate detail.

Improvement to the performance of this demonstration can be expected along the following lines:

- a. Improved modeling. More advanced modeling techniques could be investigated. Such techniques offer the possibility of modeling the terrain with fewer facets while maintaining model accuracy, or conversely, providing more detail for the same number of facets.
- b. Processor upgrades. Faster CPUs currently available would provide performance enhancements by allowing on-the-fly optimization algorithms to be used during rendering. For example, it might be possible to sort the graphics primitives by distance from the viewer for each animation frame. By drawing all primitives from back-to-front, the z-buffering hardware could be bypassed, at least partially, and the graphics hardware will work faster.
- c. Multiprocessing and parallelism. The current hardware/software configuration uses only a single process. If multiple processes could cooperate, for example, separating memory management from graphics update functions, performance will improve. Furthermore, if these processes could be distributed on multiple CPUs, even more improvements could be expected.
- d. Hardware advances. Market demand will continue to drive improvements in both CPU and graphics processor performance, probably to a quoted rate of 500,000-facets per second within 18 months.

SECTION 6

MISSION PLANNING AND REHEARSAL SIMULATION

We have applied these stereoscopic display techniques to a simulation of a mission planning and rehearsal system to demonstrate the utility of stereoscopic displays for terrain visualization. This system is useful for soliciting reactions from knowledgeable personnel and for testing new and improved techniques.

6.1 RATIONALE

The application of a 3D display presentation is most useful when knowledge of the spatial relationships among objects displayed is a required element in the decision-making process. Cases such as these arise in the air interdiction (AI) planning process at the unit level. Knowledge of enemy sensor and enemy weapons locations, with respect to terrain shape and aircraft flight profile, are key spatial elements that would benefit by a 3D presentation on an operator workstation. Additionally, in the final step, the weapons delivery process, knowledge of the spatial orientation of objects with respect to the weapons platform is critical in dispensing guided munitions. Therefore, with the establishment of those 3D cues that provide revealing stereoscopic imagery with the LCSS, this section describes the status of the development of a penetration aids display presentation for AI route planning.

6.2 SIMULATION DESCRIPTION

This simulation consists of side-by-side comparison of a 2D display capability used for route planning called CAMRA. CAMRA was developed with a stereoscopic display presentation of similar data. In effect, the stereoscopic display presentation is used to amplify and augment the operator's comprehension of the route developed on the 2D system. Use of the two systems simultaneously provides a qualitative feel for the added value of the 3D presentation. It is also possible to turn off the stereoscopic effect at any time during the simulation to show a 3D perspective view of the same scene.

The simulation is run in two sequential phases:

Phase One. 2D Display Presentation. First, the CAMRA system is demonstrated and a route is selected via this system. CAMRA, which employs a MicroVAX II/GPX workstation, shows a 2D map with political boundaries, cities, and surface-to-air missile (SAM) sites with elevation data represented by color coding. The data used in this demonstration corresponds to the Fulda Gap area of Germany.

On the first display screen, radar/aircraft intervisibility is shown simultaneously for three selectable altitudes above the terrain by an area fill technique. These visibility grids are calculated using site locations and elevation data for nearby terrain to determine screening. Visibility is shown for one SAM site at a time. The first display screen is used to orient the operator to the mission, but no operator input is involved.

In a second display (figure 6.1), the operator is presented with "relative danger contours" for an operator-selected height above the surface of the terrain. These contours are calculated based on radar screening. Contours represent points of equal danger at the given altitudes. The system can show these contours for three altitudes, one at a time. The application program will then suggest a route that minimizes the danger based on the contours. The operator can make changes to the route, altering turn points (figure 6.1). Once the operator makes changes, the program stores the route on disk for later access.

Phase Two. 3D Presentation. The Silicon Graphics (SG) 4D/70 GT stereoscopic display system is used for a mission rehearsal, "flying" the operator through the selected route. The application program, running on the SG workstation, accesses the coordinates of the selected route on the MicroVAX/GPX via an Ethernet network. A starting point on the route and altitude for the fly-through are selected by the operator. The monitor then stereoscopically displays the "view-out-the-window" along the selected route. This view-out-the-window includes terrain relief derived from DTED, overlaid with terrain coloring and a selection of objects such as buildings, smokestacks, and power pylons as derived from DFAD data. Overlaid on the terrain data is a 3D presentation of the aircraft/radar intervisibility contours using a wireframe technique.

Operator interaction for this stereo presentation is via the SG dial and button control box. This input device provides the capability to:

- a. Select a start point for the fly-through at an interesting point in the 3D volume.
- b. Stop and then resume the fly-through.
- c. Change the direction of the view at any time.
- d. Change the speed of the fly-through.




Figure 6.1. Photograph of 2D Plan View Presentation of Route Planning

6.3 INITIAL REACTION

Since the SG 4D/70 GT workstation has a practical throughput limit between 38,000 and 62,000 facets per second, we qualitatively determined an operating point that provided acceptable scene realism (complexity) and update rate for an interactive fly-through. Selected MITRE personnel viewed the mission rehearsal simulation on this workstation. A fly-through path was selected and the scene complexity and update rate were varied. The observers selected the operating point and comfort level that pleased them. An update-rate in excess of one-frame per second was satisfactory. This corresponded to a scene complexity of about 25,000 facets in each stereo field. With the selected operating point, we demonstrated the fly-through of the out-the-window stereoscopic scenes to other observers, who were favorably disposed towards the quality of the stereoscopic scene and the comfort level. We plan to perform a more complete evaluation with personnel who are experienced with MSS functions during FY90.

SECTION 7

SUMMARY AND CONCLUSION

Stereoscopic display techniques can enhance viewer comprehension of spatial relationships. We evaluated these techniques for Air Force applications, focusing on the visualization of terrain data. We developed optimized stereoscopic display techniques and implemented algorithms for the management and filtering of terrain and cultural feature databases to construct optimal geometric terrain models. We employed these techniques in constructing a simulation of a mission planning and rehearsal system to demonstrate the utility of stereoscopic 3D.

Further work remains to establish the value of stereoscopic displays for the Air Force. This work will include upgrade of the simulated mission planning and rehearsal system to provide greater scene realism and an improved frame update-rate using graphic workstations. Enhancements, such as the introduction of photographic imagery, the effects of illumination and weather, and the ability to view the attacking aircraft from the defensive site, are under consideration. We also plan to demonstrate the system to operational personnel familiar with this application in order to evaluate its effectiveness.

LIST OF REFERENCES

1. D. A. Southard, "Survey of Three Dimensional Graphics Techniques for C³I Applications," The MITRE Corporation, Bedford, MA, M89-5, February, 1989.
2. J. D. Foley, and A. Van Dam, "Fundamentals of Interactive Computer Graphics," Addison-Wesley, Reading, MA, 1984.
3. P. Bos, and T. Haven, "Field-Sequential Stereoscopic Viewing Systems Using Passive Glasses," Proceedings of the SID, Vol. 30/1, 1989, pp. 39-43.
4. D. L. MacAdam, "Stereoscopic Perceptions of Size, Shape, Distance and Direction," Journal of The SMPTE, Vol. 62, April 1954, pp. 271-293.
5. P. Tolin, "Maintaining the Three Dimensional Illusion," Information Display, Vol. 3, No. 11, December 1987, pp. 10-12.
6. L. Lipton, "Binocular Symmetries as Criteria for the Successful Transmission of Images," Proc SPIE Vol. 507, 1984, pp. 108-113.
7. N. A. Valyus, "Stereoscopy," Focal Press, London, 1966.
8. J. Baker, "Generating Images for a Time-Multiplexed Stereoscopic Computer Graphics System," Proc. SPIE, Vol. 761, 1987, pp. 44-52.
9. M. D. Abner, "Three Dimensional Visual Display for a Prototype Command and Control Work Station," Master's Thesis, Naval Postgraduate School, Monterey, CA, June 1988.
10. F. E. Harris, J. M. Yurchake, and M. J. Zyda, "Preliminary Work on the Command and Control Workstation of the Future," NPS52-88-027, Naval Postgraduate School, Monterey, CA, August 1988.
11. D. B. Smith, and D. G. Streyle, "An Inexpensive Real-Time Interactive Three Dimensional Flight Simulation System," Master's Thesis, Naval Postgraduate School, Monterey, CA, June 1987.
12. T. K. Peucker, R. J. Fowler, J. J. Little, and D. M. Mack, "Digital Representation of Three Dimensional Surfaces by Triangulated Irregular Networks (TIN)," Technical Report No. 10 (revised), Office of Naval Research, Geography Programs, Arlington, VA, 1976.
13. A. Rosenfield, and A. C. Kak, "Digital Picture Processing," Second Edition, Vol. 2, New York, Academic Press, 1982, pp. 84-96.